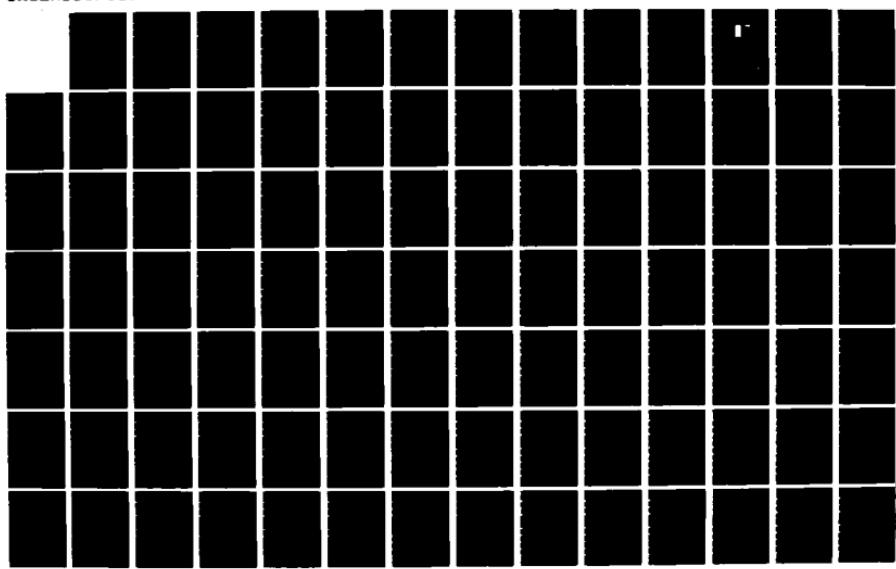
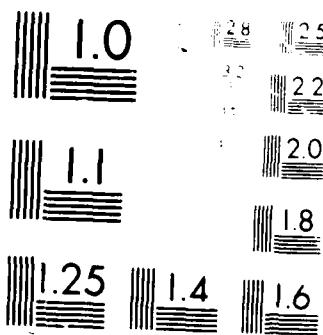


AD-A186 151 THE (S S) INVENTORY MODEL UNDER LOW DEMAND(U) AIR FORCE 1/2
INST OF TECH WRIGHT-PATTERSON AFB OH D K PETERSON 1987
AFIT/CI/NR-87-139D

UNCLASSIFIED

F/G 15/5 NL





AD-A186 151

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

DTIC FILE COPY

READ INSTRUCTIONS
BEFORE COMPLETING FORM

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFIT/CI/NR 87-139D	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) The (s,S) Inventory Model Under Low Demand		5. TYPE OF REPORT & PERIOD COVERED THESIS/DISSERTATION
7. AUTHOR(s) David K. Peterson		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS AFIT STUDENT AT: University of North Carolina		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS AFIT/NR WPAFB OH 45433-6583		12. REPORT DATE 1987
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. NUMBER OF PAGES 106
16. DISTRIBUTION STATEMENT (of this Report)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
18. SUPPLEMENTARY NOTES APPROVED FOR PUBLIC RELEASE: IAW AFR 190-1		<i>Yann E. Wolaiver</i> 260287 YANN E. WOLAIVER Dean for Research and Professional Development AFIT/NR
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) ATTACHED		

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

37 11 10 085

DAVID ALLEN PETERSON, The α, β Inventory Model Under Low Demand
Under the Direction of HARVEY M. WAGNER.

Inventory managers often rely on cost-based rules to control their stock. Unfortunately, many of these rules do not recommend inventorying infrequently demanded items due to their assumed cost efficiency. This research provides a better understanding of the impact of low-rent demands on the ASIN periodic review inventory model.

An important result from applying the (s,S) model to newsvendor demand is that negative reorder points are commonly recommended. In other words, the firm does not reorder stock until a sufficiently large number of backorders are "on-the-books." For this cost cannot tolerate lengthy delays in ordering, the (s,S) model may be constrained to allow only nonnegative reorder points. While maintaining reorder points increases the inventory service, it also increases total costs. Fortunately, the (s,S) model is relatively insensitive to misspecifying the mean demand per period.

We must show that approximately optimal (s, S) inventory policies are feasible alternatives to computationally burdensome (s, S, M) policies under low demand. In fact, there is very little cost degradation from using approximately optimal (s, S) policies rather than optimal (s, S, M) policies when demand is Poisson distributed. This is true whether the demand parameters are known or estimated from historical demand data.

Finally, we demonstrate that useful approximations for (s, S) control operating characteristics under low mean demand are easily obtained. These approximations allow inventory practitioners to quickly evaluate the impact of business decisions on their operating costs and service.



THE (s,S) INVENTORY MODEL UNDER
LOW DEMAND

by

David K. Peterson

A Dissertation submitted to the faculty of
The University of North Carolina at Chapel Hill
in partial fulfillment of the requirements for the
degree of Doctor of Philosophy in the
School of Business Administration.

Chapel Hill

1987

Approved by:

Henry W. Johnson
Adviser

Dana Quade
Reader

Michael Elbawatt
Reader

Ann E. Mancheck
Carney G. Fisher

DAVID KENT PETERSON. The (s, S) Inventory Model Under Low Demand (Under the direction of HARVEY M. WAGNER).

Inventory managers often rely on cost-based rules to control their stock. Unfortunately, many of these rules do not recommend inventorying infrequently demanded items due to their assumed cost structures. This is especially troublesome for a business where the lack of some infrequently used items can interrupt the firm's smooth operation. This research provides a better understanding of the impact of low mean demand on the (s, S) periodic review inventory model. We assume a fixed delivery lag, complete backlogging of unfilled demand, a fixed ordering cost, and linear purchasing, holding and shortage penalty costs.

We examine the impact of low mean demand per period on the (s, S) periodic review inventory model. An unexpected result from applying the (s, S) model to low mean demand is that negative reorder points are commonly recommended. In other words, the firm does not reorder stock until a designated number of backorders are "on-the-books." For firms that cannot tolerate lengthy delays in ordering infrequently demanded but critical parts, the (s, S) model may be constrained to allow only nonnegative reorder points. While constraining reorder points increases the inventory service, it also increases total costs. Fortunately, the (s, S) model is relatively insensitive to misspecifying the mean demand per period. That is, using an estimated mean demand slightly higher or lower than the actual value does not increase total costs appreciably.

We also show that approximately optimal (s, S) inventory policies are feasible alternatives to computationally burdensome optimal (s, S) policies under low demand. In fact, there is very little cost degradation from using approximately optimal (s, S) policies rather than

optimal (s, S) policies when demand is Poisson distributed. This is true whether the demand parameters are known or estimated from historical demand data.

Finally, we demonstrate that useful approximations for (s, S) policy operating characteristics under low mean demand are easily developed. These approximations allow inventory practitioners to readily evaluate the impact of business decisions on their inventory costs and service.

Acknowledgements

I wish to sincerely thank my adviser, Dr. Harvey M. Wagner, for his guidance and encouragement. I would also like to thank the other members of my doctoral committee for their many valuable suggestions and contributions to this research: Dr. Richard Ehrhardt, Dr. Nancy Hyer, Dr. Ann Maruchek, and Dr. Dana Quade. My thanks also go to the U. S. Air Force for providing me the opportunity to earn my doctorate, and to Gean McBane who typed the dissertation. Finally, I wish to thank my wife, Kay, for her unfailing support and patience during the past three years. To her goes much of the credit for my successfully completing this doctoral program.

TABLE OF CONTENTS

I. Introduction and Literature Review	1
II. A Comparison of Optimal and Constrained Optimal (s,S) Inventory Policies Under Low Demand (Known Mean)	23
III. A Comparison of Constrained Approximately Optimal (s,S) Inventory Policies Under Low Demand (Known Mean)	39
IV. An Evaluation of Constrained Approximately Optimal (s,S) Inventory Policies Under Low Demand (Estimated Mean)	53
V. Regression Derived Approximations for Constrained Optimal and Approximately Optimal (s,S) Inventory Policy Operating Characteristics	64
VI. Conclusion	77
Appendix I: 480 Optimal (s,S) Policies	82
Appendix II: Supplementary Simulation Results	93
Appendix III: Operating Characteristic Approximation Models for the Constrained Optimal and Power Approximation Policies	95
References	98

CHAPTER 1

INTRODUCTION AND LITERATURE REVIEW

1.1 Low Demand Inventory Management

Modern inventory managers often rely on cost-based rules to control their stock. For example, the U. S. Air Force stockage model for expendable items [44] and the economic order quantity inventory model [47] both rely on "economic tests" to determine if an item should be stocked. Unfortunately, many of these tests do not recommend stocking infrequently demanded items due to their assumed cost structures. This is especially troublesome for a business where the lack of some infrequently used items can interrupt the firm's smooth operation. Surprisingly, although low demand items pose significant practical problems, researchers have made limited progress on this inventory management issue.

1.1.1 Low Demand Items - A Military Example

Today's sophisticated weapon systems and restrictive budgetary constraints have placed a premium on effective inventory management within the military. Two critical decisions facing military inventory managers are: 1) should an item with an historically low demand rate be stocked, and 2) if stocked, what is the appropriate quantity?

These questions are not trivial when the lack of a single critical item can ground a multi-million dollar aircraft. Appropriate answers to these questions assume an even greater importance when weapon systems must operate from locations supported by lengthy and complex logistics systems.

1.1.2 Demand Pattern

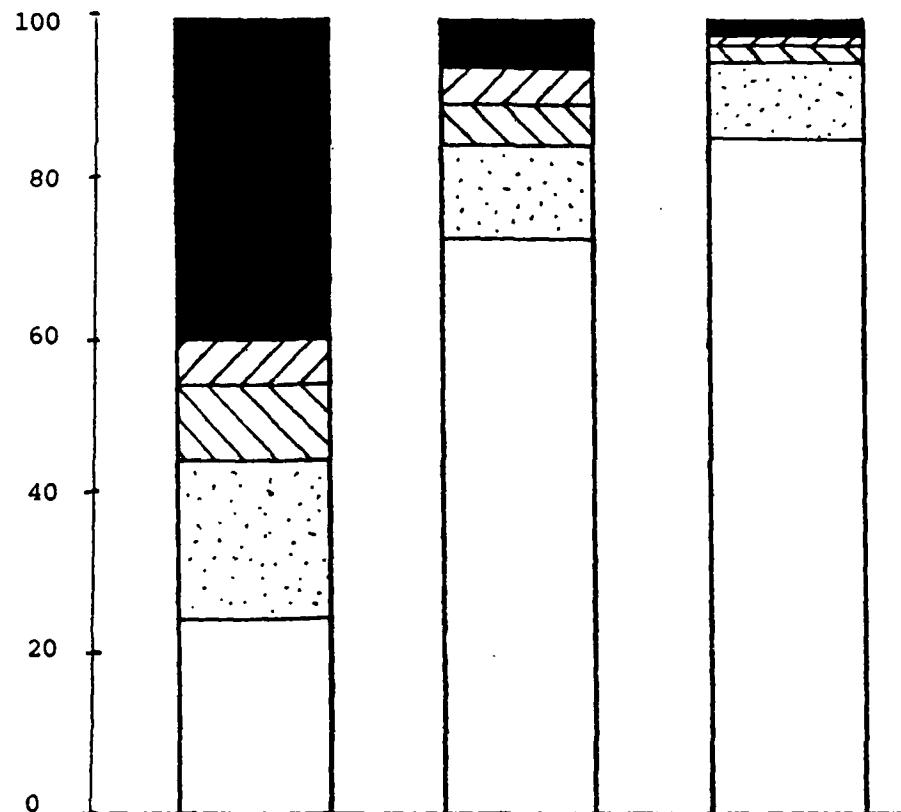
Over the past three decades, a consistent pattern has emerged concerning parts demand for military aircraft, ships and vehicles. The studies cited below show that the vast majority of items comprising a weapon system are rarely demanded.

In 1956, Brown [5] reported that of over 15,000 unique parts applicable to the B-47 aircraft, 25% experienced no demands at a base supporting approximately 50 aircraft over a year's time frame. In 1963, Feeney et al [19] found that 74% of the high value, repairable spares stocked at Andrews Air Force Base were not demanded during a six month period. Haber and Sitgreaves [22], in 1970, reported that from over 25,000 potential repair parts, fewer than 20% were demanded during 61 submarine patrols. This pattern also holds for land vehicles, as Frazza [20] discovered in 1982 after aggregating two years of parts demands for 107 5-ton tractor trucks.

Figure 1.1 summarizes several of these studies. Although the part populations and time periods are not strictly equivalent, the studies' overall demand patterns are similar. In essence, Figure 1.1 demonstrates that managing low demand items is an important challenge to military inventory managers.

In studying low demand items, inventory researchers have frequently found the Poisson probability distribution a useful representation of parts demands in military systems. Brown [5] examined a sample of 100 B-47 airframe parts having sufficient demand history to judge their probability distribution. She concluded that 79 "seemed to have demand patterns over time that can be described by the Poisson distribution [5]." Feeney et al [19], in evaluating a new base stockage policy, used

Figure 1.1
Observed Demand Frequencies



Brown
1956
[12 months] Feeney
1963
[6 months] Haber
1970
[61 patrols]

Demand Studies
[Time Frame]

Item Demands
0 1 2 3 4+

the Poisson distribution to represent the "low variability case ... the sort of variability we would expect if we were dealing with a random failure process [19]." Haber and Sitgreaves [22,23] used the Poisson distribution to represent parts demand during submarine patrols. Frazza [20] tested the Poisson failure assumption made by the U. S. Army's Selected Essential Item Stockage for Availability Method (SESAME) model. She found that for weapon systems "with a large number of items and short usage duration we cannot reject the Poisson distribution [20]." She also reached this same conclusion for trucks having more extensive usage (over 60,000 miles).

Several researchers have assumed a Poisson distribution for low demand items in commercial (non-military) environments. Whitin and Youngs [49], in 1955, developed a method for calculating optimal inventory levels for a fleet of taxicabs assuming Poisson demand. In 1956, Heyvaert and Hurt [24] developed a method for determining optimal inventory levels assuming Poisson demand. Schaefer [41], in 1983, modeled an inventory system for a maintenance center experiencing Poisson demand. Finally, studies such as Davis [6] (in 1952) have actually examined data for commercial aircraft spares and determined that the hypothesis of Poisson distributed demand could not be rejected. Thus, many inventory researchers have assumed Poisson distributed demand for studying military and commercial low demand items. This assumption has been supported by several evaluations of actual usage data, and, as Davis [6] states:

discrepancies between the theory and data (though they may be significant in a statistical sense) are small enough that the exponential theory [Poisson demand] may be regarded as a useful approximation of certain classes of actual failure distributions.

The Poisson distribution has several other characteristics that make it attractive in modeling inventory systems. First, its mathematical tractability facilitates adapting queueing theory to inventory models. Second, if the time between arrivals to a system is random (thus having a negative exponential distribution), the number of arrivals in any specified time period is Poisson distributed. If each arrival demands only one unit, then demand over any specified time interval is also Poisson distributed. Finally, the simple Poisson distribution is specified by a single parameter, the mean (which also equals the variance), thus simplifying the task of estimating the demand distribution.

While these are attractive theoretical features, one criticism of using the simple Poisson distribution is that variances estimated from empirical demands are usually much larger than the estimated means. Nevertheless, as discussed above, the Poisson distribution plays a key role in the low demand inventory literature.

1.1.3 The Inventory Model

This research study assumes that an (s, S) policy is used for the replenishment rule. At the start of each period t , the stockage position X_t , which is stock on hand and on order, is reviewed. A replenishment order is placed whenever X_t falls below the reorder point s ; the order is received after a fixed lead time of L periods, prior to that period's demand. The size of the order x_t is chosen so that Y_t , the stockage position after ordering, is brought up to the stockage objective S .

Mathematically, the (s, S) policy operates in the following manner when all unfilled demand is backordered [45, 47]:

$$(1.1) \quad Y_t = \begin{cases} i_t + x_{t-L} + \dots + x_t = s & \text{if } x_t < s. \\ x_t & \text{otherwise,} \end{cases}$$

where i_t is inventory on hand at the start of period t , and x_v is the amount ordered in period v .

The model assumes an infinite planning horizon with each period's demand q_t being independently and identically distributed. If q_t exceeds the on hand inventory i_t , the deficit is backordered ($q_t - i_t$) and filled from future receipts. Hence $X_{t+1} = Y_t - q_t$.

Replenishment cost $c(x)$ is given by

$$c(x) = \begin{cases} 0 & \text{if } x = 0 \\ K + cx & \text{if } x > 0, \end{cases}$$

where x is the amount ordered. The per period holding cost and shortage penalty cost are linear functions

$$l(j) = \begin{cases} h_j & \text{if } j \geq 0 \\ p(-j) & \text{if } j < 0, \end{cases}$$

where $j = i_t - q_t$.

Finally, the (s, S) policy is found by minimizing the long-run average (non-discounted) cost per period [47]. The relevant expression for the objective function is

$$(1.2) \quad K/(1+M(D)) + \sum_{z=0}^D L(S-z)r(S-z),$$

where the functions $M(\cdot)$, $L(\cdot)$, and $r(\cdot)$ are defined as follows. The function $M(\cdot)$ is the quantity

$$(1.3) \quad M(D) = \sum_{t=1}^{\infty} p^t(D),$$

where P^t is the cumulative distribution of the t -fold convolution of $p(q)$ (the probability of q demands in any period) and $D=(S-s)$. The function $L(\cdot)$ is the expected holding and penalty cost function

$$(1.4) \quad L(y) = \begin{cases} \sum_{Q=0}^y h^*(y-Q)p^{L+1}(Q) + \sum_{Q=y+1}^{\infty} p^*(Q-y)p^{L+1}(Q) & \text{if } y \geq 0 \\ \sum_{Q=0}^{\infty} p^*(Q-y)p^{L+1}(Q) & \text{if } y < 0, \end{cases}$$

where Q is the cumulative demand over L periods ($q_{t_1} + \dots + q_{t+L-1}$) and $p^t(q)$ is the t -fold convolution of $p(q)$. The function $r(\cdot)$ is the normalized renewal quantity

$$(1.5) \quad r(S-z) = (\sum_{t=1}^{\infty} p^t(z)) / (1+M(D)) \quad \text{for } 0 \leq z \leq D.$$

Since all demands are eventually filled, the expected purchase cost per period, which depends on c , is a constant and thus unaffected by (s, S) .

1.2 Purpose

This research extends knowledge about the (s, S) periodic review inventory model's behavior to the case of low mean demand. Specifically, the research addresses:

1. Operating Characteristics. How does the (s, S) inventory model behave for low mean demand values? How much are costs increased over an optimal policy when inventory policies are constrained to nonnegative reorder points? Also, how are costs increased when constrained inventory policies are used with misspecified demand means?

2. Reorder Policy Approximations. Is there an easy-to-compute approximately optimal (s, S) inventory policy that performs well when mean demand is low and known, and when the reorder point is constrained to be nonnegative? By how much do costs increase when mean demand must be estimated from historical data?

3. Performance Approximations. Can easy-to-compute approximating formulas for (s, S) inventory policy operating characteristics be developed? How can inventory managers use such approximations to assess the expected performance of (s, S) inventory policies when mean demand is low?

1.3 Literature Review

During the past 35 years, the (s, S) periodic review inventory model has evolved into a powerful inventory management technique. Sections 1.3.1 through 1.3.4 below present the research evolution of this model. The model's response to parameter changes is reviewed in Section 1.3.5, and demand estimation techniques are discussed in Section 1.3.6. Finally, in Section 1.3.7, conclusions are drawn from the research literature and several research questions to be addressed in this study are proposed.

Research on the (s, S) periodic review inventory model has passed through four distinct phases (Figure 1.2). These phases are: 1) early dynamic models, 2) proving the optimality of (s, S) policies, 3) devising efficient computational algorithms, and 4) searching for approximation methods that are effective when demand parameters must be estimated.

Figure 1.2

The Evolution of (s,S) Policy Research

Year	Early Dynamic Models	Proving (s,S) Optimality	Efficient Algorithms	Approximation Methods
1950				
1	-Arrow, Harris, & Marschak			
2	-Dvoretzky, Kiefer			
3	& Wolfowitz			
4	-Bellman, Glicksberg, & Gross			
5				
6				
7				
8				
9				
60		-Scarf		
1				
2				-Roberts
3		-Iglehart		
4				
5			-Veinott	-Wagner, & Wagner
6				O'Hagan, & Lundh
7				
8				-Johnson
9				
70				
1				
2				
3				
4				
5				-Naddor
6				-Ehrhardt
7				
8				
9				
80				-Freeland & Porteus
1				
2				
3				-Ehrhardt &
4				Mosier
5				-Ehrhardt/ Porteus
6				
7				

1.3.1 Early Dynamic Models

Arrow, Harris and Marschak [1] in 1951 published the earliest (s, S) periodic review inventory model. Their dynamic model separates the inventory replenishment strategy into two parts - the "current decision and all subsequent decisions [40]." They treat the unbounded horizon, stationary case. Their model selects two inventory control parameters, S (the stockage objective) and s (the reorder point). These parameters are chosen to minimize expected loss (costs) per period. The model represents customer demand probabilistically, and allows reordering stock at fixed time intervals. They derived equation (1.2) for a continuous demand density.

In 1952, Dvoretzky, Kiefer and Wolfowitz [7] developed a generalized solution to the inventory problem using dynamic programming. They prove a solution exists to a functional equation defining an optimal inventory policy [2,40]; the solution, however, need not yield an (s, S) policy [7]. Specifically, their extremal equation is

$$(1.7) \quad f(x) = \inf_{y>x} \{ V(x,y) + \alpha f(0)[1-P(y)] + \alpha \int_0^y f(y-q)dP(q) \}$$

where $V(x,y)$ is the expected current period costs and α , $0 < \alpha < 1$, is the discount rate. Note that they treat the case where any unsatisfied demand is lost and not backordered, and lead time is zero.

In 1953, Dvoretzky, Kiefer and Wolfowitz [9] studied when an (s, S) policy is optimal given discrete (integer valued) demand, and "give some sufficient conditions for establishing that the optimal policy is an (s, S) policy for the single-stage inventory problem [2]." They

assume a fixed stockout penalty cost, and an ordering cost comprised of fixed and linear components [2].

The last of the early models is by Bellman, Glicksberg and Gross [3] in 1955, who built upon the preceding dynamic inventory models. Theirs is the "first analysis of the inventory equation whose purpose was to determine the form of an optimal policy [40]." They show [40]:

the optimal policy is defined by a sequence of critical numbers ... if the stock level at the beginning of period n is below x_n , an order for the difference is placed. No ordering is done during this period if the stock level exceeds x_n .

They employ a dynamic version of (1.7) for a finite horizon; specifically, they use

$$(1.8) \quad f_n(x) = \min_{y \geq x} [k(y-x) + \int_y^\infty p^*(s-y)p(s)ds + f_{n-1}(0) \int_y^\infty p(s)ds + \int_0^y f_{n-1}(y-s)p(s)ds].$$

Note that the $k(y-x)$ function in their model assumes no set-up or holding costs [3], and lead time is zero.

These early research efforts defined the (s, S) periodic review inventory model's basic characteristics and paved the way for establishing the optimality of (s, S) policies.

1.3.2 (s, S) Policy Optimality

In 1960, Scarf [39] published the first proof of (s, S) policy optimality for a finite horizon. His inductive proof shows, with certain assumptions, that the form of an optimal inventory policy for each period n is (s_n, S_n) . The proof establishes that the cost function is K -convex, and that an (s_n, S_n) policy minimizes expected costs over

an n period planning horizon. Specifically, Scarf treated the functional equation

$$(1.9) \quad f_n(x) = \min_{y \geq x} \{c(y-x) + L(y) + \alpha \int_0^{\infty} f_{n-1}(y-q)p(q)dq\},$$

where

$$(1.10) \quad c(z) = \begin{cases} 0 & \text{if } z=0, \\ K + c^*z & \text{if } z>0. \end{cases}$$

His method of proof extends immediately to the case of a positive lead time with complete backlogging.

Iglehart [25], in 1963, proved an (s, S) policy's optimality over an infinite planning horizon [40]. He proves an (s, S) policy's optimality under both a total discounted cost criterion and an average (non-discounted) cost per period criterion [43, 25, 40]. The starting point in his analysis is (1.9); he examines the limiting behavior of f as $n \rightarrow \infty$. Having established the (s, S) model's basic characteristics and proved its optimality, the stage was set for developing efficient computational procedures.

1.3.3 Efficient (s, S) Computations

In 1965, Veinott and Wagner [45] developed the first efficient method to compute optimal (s, S) policies. The model assumes an infinite planning horizon, integer-valued demand, and complete backlogging of unfilled demand. Their algorithm is comprised of a three-step procedure that seeks to minimize expected total costs per period (1.2) or expected discounted costs [47]:

Step 1. Calculate lower and upper bounds on both s and S , and thereby determine lower and upper bounds for an optimal $D \in S-s$.

Step 2. For each value of D between the upper and lower bounds found in Step 1, compute the stationary distribution $r(v)$, and find a corresponding optimal S by minimizing D

$$\sum_{z=0} L(S-z)r(S-z).$$

Step 3. For each of the pairs of D and S obtained in Step 2, calculate the associated value of [equation 1.2] and select a pair that gives the minimum for [equation 1.2].

In contrast with the Veinott-Wagner algorithm, Johnson [27] in 1968 and Federgruen and Zipkin [18] in 1984 suggest policy iteration methods. Their procedures begin with an approximately optimal (s, S) policy, and iteratively search for policies that further minimize long-run average costs. They claim their techniques more efficiently compute an optimal (s, S) policy. Certainly, they benefit from using an approximately optimal (s, S) policy as a good starting point.

1.3.4 Approximation Methods

Approximately optimal policies are an attempt to avoid the computational burden and demand information requirements of optimal (s, S) policies. Roberts [36], in 1962, used renewal theory to characterize the limiting behavior of an optimal (s, S) policy as a function of the model parameters. His results for s and $D = (S-s)$ are [36]

$$(1.11) D = ((2K\mu/h)^5) + o(D)$$

$$(1.12) \int_s^{\infty} (x-s) dP^{L+1}(q) = D/(1 + p/h) + o(D),$$

where $o(D)$ is the approximation error, and " $o(D^*)/D^*$ converges to zero as D^* becomes infinite [11]."

In 1965, Wagner, O'Hagan and Lundh [46] compared exact and approximately optimal policies in an empirical (computer simulation)

study. For the Poisson distribution, they show even the best performing approximations (Roberts and Normal) deteriorate as mean demand per period increases, and do not approximate the difference $D = (S-s)$ well.

In 1975, Naddor [34] presented rules to find approximately optimal (s,S) policies using the Normal Approximation. Porteus [35] calls Naddor's approximation the "Analogy" method "because it selects s and S by using approximately optimal policies for analogous models [35]." The details of this method are described in Chapter 3 below.

Ehrhardt [10], in 1976, continued research into approximating optimal policies and derived the Power Approximation. "The Power Approximation is an algorithm for computing approximately optimal values for (s,S) using only the mean and variance of demand [10]." By computing a value for D and s , and thus $S = D+s$, the policy parameters are specified.

The equations for determining D and s were found by fitting the values from optimal (s,S) policies in a regression model comprised of the model's parameters. For Poisson distributed demand, Ehrhardt reports the Power Approximation outperforms other methods such as the Normal Approximation [10]. Ehrhardt and Mosier [15], in 1984, revised the Power Approximation to: 1) improve the D value calculation, and 2) ensure system parameters are homogenous in the units of demand measurement. The computational details of their Revised Power Approximation are also described in Chapter 3 below. Ehrhardt and Mosier evaluated the Revised Power Approximation in a factorial design with 288 cases. Using deviation from optimal total cost as the criterion, they found 99.7% of the cases were within 5% of optimum.

In 1976, Klincewicz [31] examined the Power Approximation's performance when there is a 25% chance of zero demands in any period. A compound negative binomial distribution, with a variance-to-mean ratio of 9, is used to represent demands. Klincewicz found that the "relative performance of the Power Approximation and the Normal Approximation is a function of mean demand: for items with low demand ($\mu = 2,4$), the Power Approximation gives a lower average total cost, whereas for items with high mean ($\mu = 8,16$), the Normal Approximation is preferred [31]."

In 1980, Freeland and Porteus [21] empirically tested "a new method ... for obtaining approximately optimal (s,S) inventory policies ... [by] approximating stochastic dynamic programs [21]." The new method performs slightly better than the empirical approximation of Wagner, O'Hagan and Lundh [46] when demand follows either a normal or negative binomial distribution. Unfortunately, the new method is more computationally burdensome than the empirical method. In 1985, Porteus [35] empirically tested 17 different approximation methods. He concludes that [35]:

depending on the range of values that apply in a given practical situation, either (i) any of a large number of methods will yield good performance, or (ii) a carefully selected method can achieve superior performance.

Finally, in 1985, Ehrhardt [16] reported a set of approximations for (s,S) inventory system operating characteristics. These approximations provide inventory managers relatively quick and accurate estimates of an (s,S) inventory system's performance. The five approximations are for the expected: 1) total cost per period, 2) holding cost per period, 3) replenishment cost per period, 4) backlog cost per period, and 5) backlog frequency. Fit to a factorial design

with 288 cases, the average percent deviation from actual values varied for each approximation; specifically, total cost (2.8%), holding cost (3.9%), replenishment cost (1.4%), backlog cost (6.2%), and backlog frequency (0.7%).

1.3.5 Impact of Parameter Variations on the (s,S) Model

There are five main parameters that define the (s,S) periodic review inventory model used in this study. These parameters are: 1) ordering setup cost K , 2) unit holding cost h , 3) unit penalty cost p , 4) delivery lag L , and 5) mean demand μ . As these parameters change, the (s,S) policy parameters (s , S , D) change in a determinable manner. Table 1.1 summarizes the changes in s , S , and D when h is normalized to 1 [47]. Note S need not increase with p or μ , and D need not increase with μ [47].

Table 1.1
Parametric Impacts on (s,S) Inventory Policy System Parameters

Parameter	s	S	D
As: K ^	v	^	^
p ^	^	-	v
L ^	^	^	^
μ ^	^	-	-

While many researchers have studied the impact of parameters and (s,S) policies, little attention has been devoted to truly low mean demand values. Table 1.2 summarizes the mean demand values used in several (s,S) policy studies. These mean demand values are however,

much higher than commonly experienced by many military weapons systems (Figure 1.1). Further research is needed to evaluate the impact of low mean demand on the (s, S) periodic review inventory model.

Table 1.2

Previously Studied Values of Mean Demand per Period

Study	μ Values
Wagner, O'Hagan and Lundh	1, 2, 3, 4 ...
Ehrhardt	2, 4, 8, 16
Klincewicz	2, 4, 8, 16
Freeland and Porteus	2, 4, 8, 16 ...
Federgruen and Zipkin	2, 6, 20, 60
Porteus	2, 6, 20, 60

1.3.6 Demand Estimation

Over the past 35 years, the (s, S) periodic review inventory model has evolved into a well-accepted and useful inventory management tool. It has a sound theoretical basis, and proven optimality under certain operating assumptions. Further, optimal (s, S) policies may be computed either exactly or approximately as the situation warrants. Unfortunately, although some parametric impacts are well known, inventory research has not studied the impact upon (s, S) inventory policies of truly low mean demand, such as less than one demand per period.

As Wagner states in his 1980 inventory research portfolio: "What are effective replenishment formulas when observed demand is mostly zero valued (for example, 40 weeks of no demand interspersed among 12

weeks of positive demand)? [48]" Silver in his 1981 review of inventory research makes the same point [42]:

Most models that are usable ... assume unit-sized demand transactions or relatively smooth demand patterns. These assumptions are inappropriate under conditions of ... intermittent demand.

Silver further states, "It would be very helpful to have practical means ... of setting appropriate values of the control parameters."

The (s,S) inventory policy literature has treated the statistical problem of parameter estimation in two distinct manners. Dvoretzky, Kiefer, and Wolfowitz [8] and Scarf [37,38] assume prior distributions about the demand parameters and attempt to identify an optimal statistical estimation procedure. The majority of (s,S) policy research, however, has used standard statistical procedures such as the sample mean and variance or exponential smoothing to estimate demand parameters. For example, Klincewicz [31], Ehrhardt [10,11], Blazer [4], or Jacobs [26].

Unfortunately, as Scarf [40] notes, the attempts "to introduce a priori distributions that are successively refined as more and more demand data become available ... [have] not yet proved to be particularly fruitful." Thus, researchers have done the next best thing, and compared the performance of (s,S) policies using statistically estimated demand parameters to the performance when the demand parameters are known. This approach provides a bound on the (s,S) policy performance if an optimal statistical demand estimator could be derived. As of this writing, finding an optimal statistical estimation procedure for the model studied in this research remains a difficult and open research issue.

Much of the research on the Power Approximation has employed the method of comparing (s, S) policy cost performance with known parameters to that with statistically-estimated demand parameters. In 1976, Ehrhardt [10] examined the Power Approximation's performance by simulating a 72 item inventory system, operating for 5200 periods, where the (s, S) policy is revised every 26 periods with updated demand parameters. The demand parameters are updated using the sample mean and sample variance for the previous 26 periods' demand. When demand is Negative Binomial and the variance-to-mean ratio is 9, Ehrhardt found that using statistically-estimated demand parameters increases total system costs 11.5% above optimal.

Klincewicz [30], in 1976, examined the Power Approximation's performance when a biased variance estimate replaces the sample variance. Using the same inventory simulation as Ehrhardt [10], Klincewicz' study inflated or deflated the sample variance by a constant factor β , where β minimizes an item's expected total costs per period. For a 12 item system, with Negative Binomial demand (variance/mean = 9), Klincewicz found that Power Approximation policies using biased variance estimates showed only a small improvement in total costs over policies using unbiased estimates. This improvement ranged from 0.1% to 5.6%, with items having low mean demand and high penalty costs benefitting the most.

Jacobs [26], in 1985, searched for better statistical methods to estimate demand parameters than the traditional sample mean and variance. He evaluated several intentionally-biased, statistically estimated demand parameters using the same inventory simulation and 72 item system as Ehrhardt [10]. For estimating the mean, Jacobs compared

the sample mean with an exponentially-smoothed average. For estimating the variance, he compared the sample variance to an exponentially-smoothed mean absolute deviation (MAD) biased with several different factors. The bias factors were regression-derived by fitting model parameters to total cost values. Jacobs found that using the exponentially-smoothed average and biased, exponentially-smoothed MAD yielded deviations in system costs around 7.5% and 8.5% above the full information case for the Power and Analogy Approximations, respectively. This compares with a 12.5% and 13.8% degradation for the Power and Analogy Approximations, respectively, operating under the sample mean and variance. Thus, Jacobs recommends using the exponentially-smoothed average and the biased, exponentially-smoothed MAD for inventory systems operating under either the Power or Analogy Approximation.

1.3.7 Summary

This literature review has shown there is a need for research into the impact of low mean demand upon the (s,S) periodic review inventory model. Over the past 35 years, researchers have refined this model into a powerful inventory management technique. Unfortunately, most prior research assumed mean demand values much larger than experienced by many critical yet infrequently demanded items in the business or military sectors. Thus, although low demand items pose significant practical problems, researchers can offer only limited insight into how the (s,S) periodic review model performs in a low demand environment. This research provides a better understanding of the impact of low mean demand on the (s,S) periodic review inventory model.

1.4 Contributions of this Research

The major contributions of this research are as follows:

1. To explain how an optimal (s, S) inventory policy frequently recommends a negative reorder point under conditions of low mean demand. To demonstrate that constraining reorder points to nonnegative values raises total system costs appreciably. And to show that the constrained (s, S) periodic review inventory model is not overly sensitive to misspecification of the mean demand per period when demand is low.
2. To show that both the Power Approximation and the Analogy Approximation perform well when mean demand is low, even when mean demand must be estimated from historical data.
3. To provide useful approximations for the operating characteristics of constrained (s, S) inventory policies when mean demand is low. These approximations allow easy computation of: 1) total system cost, 2) holding cost, 3) replenishment cost, and 4) backlog protection.

1.5 Overview of the Research

The remainder of this study examines the interaction of low mean demand and (s, S) inventory policies in greater detail. Chapter 2 presents the impact of low mean demand on the optimal (s, S) periodic review inventory model. It also examines the impact of constraining the reorder points to nonnegative values. Finally, Chapter 2 makes a

preliminary test of how sensitive (s,S) policies are to accurately estimating the mean demand parameter.

In Chapter 3, the performance of two constrained, approximately optimal (s,S) inventory policies is examined when the mean demand is low but known. Further, Chapter 3 examines one of the approximation's performance under an interpolated and extrapolated parameter set. Chapter 4 continues the comparison of the two approximation policies and examines their performance in a statistical demand environment.

Chapter 5 develops regression derived operating characteristic approximations for the constrained optimal and constrained Power Approximation inventory policies. Several of these approximations are applied to practical inventory management problems. Finally, Chapter 6 summarizes the findings of this research, and presents several implementation issues. Chapter 6 also describes areas for future research.

CHAPTER 2

A COMPARISON OF OPTIMAL AND CONSTRAINED OPTIMAL (s, S) INVENTORY POLICIES UNDER LOW DEMAND (KNOWN MEAN)

2.1 Introduction

Prior research has not studied the impact of low mean demand on the (s, S) periodic review inventory model. Nor has research examined how constraining the reorder points to be nonnegative affects the performance of (s, S) policies. Finally, little is known about the cost sensitivity of misspecifying mean demand.

This chapter examines all of these issues. Under low mean demand, (s, S) inventory policies may have negative reorder points. In this study, 38% of the optimal policies examined had negative reorder points. A negative reorder point implies that the firm postpones reordering until a backlog of demands occurs. Constraining reorder points to nonnegative values raises the system's average total costs while providing better service. The analyses below establish that the constrained (s, S) policies are not overly sensitive to misspecifying mean demand.

Section 2.2 describes the inventory model used by this research, including parameter specification and performance measurement criteria. In Section 2.3 the impact of low mean demand on optimal (s, S) policies is discussed. Section 2.4 shows the impact of constraining optimal (s, S) policies to nonnegative reorder points under low demand. Finally, Section 2.5 examines the cost penalty of misspecifying statistical

demand parameters. Section 2.6 then draws conclusions from these results.

2.2 Methodology

2.2.1 Inventory Model Assumptions

The (s, S) periodic review inventory model is used for all three research issues presented below. This model makes the following assumptions [45], which were detailed in Chapter 1:

1. The replenishment policy is for a single product stocked at a single location.
2. The planning horizon is unbounded.
3. The random amount demanded in successive periods is independently and identically distributed.
4. The lead time between placing and receiving an order is constant.
5. Future costs are not discounted.
6. The setup cost for placing an order is fixed.
7. The purchase cost varies linearly with the number of units ordered.
8. Any demand not satisfied in a period is backordered until it is ultimately satisfied.
9. The per period holding and shortage penalty costs are linear in the amount of inventory on hand or backordered at the end of the period, respectively.
10. Inventory on hand is conserved. There is no loss, spoilage, or condemnation.
11. The objective is to select a policy that minimizes the long-run average cost per period.

These assumptions imply that there is an optimal policy of the (s, S) type. The mathematical description of this model was given in Section 1.1.3.

2.2.2 Policy Operating Rules

The (s, S) periodic review inventory model follows a specific operating rule sequence [45, 47]. Inventory on hand and on order is reviewed at the beginning of each period. An order is placed for a positive, integral amount whenever the amount on hand plus on order is less than the reorder point s . The amount ordered is the difference between the stockage objective S and the amount on hand and on order.

All deliveries are received the lead time number of periods later, after any ordering but prior to any demands occurring in that period. If any period's demand exceeds the amount of stock on hand, the shortfall is backordered until filled by future receipts.

2.2.3 Policy Implementation

The Veinott-Wagner algorithm for optimal (s, S) policies, described in Chapter 1, was implemented in a computer program by Kaufman [28] in 1976. His computer program offers three options. First, it can compute optimal (s, S) policies. Second, it can compute operating characteristics for any researcher-specified (s, S) policy. Finally, it can calculate the best (s, S) policy for any researcher-specified $D = (S-s)$ value. Describing the program's operation, Kaufman says [28]:

The Veinott-Wagner algorithm provides a finite range for D and S in which an optimal inventory policy (s^*, S^*) must lie. Using the fact that ... [total expected cost per period] is convex in S for any given D , Program I finds the minimizing value of S for each D in the search range. The pair (S, D) with the smallest total expected cost per period is the optimal policy.

2.2.4 Parameter Set

Table 2.1 is a compendium of the 480 parameter cases initially evaluated in this study. The ordering cost and holding cost values are similar to those used by U. S. Air Force inventory policies [44]. They represent the fixed charge for submitting an order for additional materials, and the charge for holding one unit of inventory for one period, respectively. The two penalty cost levels correspond to 80% and 90% service levels, and represent the percent of time that an optimally-controlled system incurs no backlogs [47]. These service levels bracket the current Air Force 86% target. The lead times common to many earlier (s, S) inventory policy studies, are the number of time periods that elapse between placing and receiving an order. Finally, the range of mean demand represents values lower than previously investigated (Table 1.2) yet commonly experienced with military systems (Figure 1.1).

Table 2.1
Experiment with 480 Parameter Cases

Parameter	Values
K	\$3, 20
h	\$.1 (.2) 0.7
p	4*h and 9*h
L	0, 2, 4
μ	0.1 (.1) 1.0 (Poisson)

These parameters imply, using as a rough approximation an economic order quantity inventory model analogous to the (s,S) periodic review inventory model [47], that the shortest time between orders will be 3 periods and the longest time 70 periods. In other words, there will be from less than one to seventeen orders per year. This time horizon and reorder frequency is very realistic from a practitioner's viewpoint. A more precise description of the 480 cases is given in Section 2.3.

2.2.5 Explanation of Multi-Item System

The 480 parameter cases and their corresponding (s,S) policies may be viewed in two distinct ways. First, each parameter case, and corresponding (s,S) policy, represents a result that may be examined individually. This "casewise" view is especially useful when determining the impact of parametric variations. Another view, however, is the aggregation of the 480 cases into a hypothetical inventory system. This "system" view is appealing to inventory managers, who commonly collect aggregate inventory performance statistics. Whenever the system view is used in this research, the results are clearly identified as either "aggregate" or "system" statistics.

2.2.6 Performance Measurement Criteria

The criterion that measures the degradation in performance between optimal and constrained optimal (s,S) inventory policies is the percentage by which a given operating characteristic for the constrained optimal (co) policy exceeds its optimal (o) policy counterpart. This measure is computed as

$$(2.1) \quad d_c = ((C_{co} - C_o)/C_o) \times 100$$

where

C_{CO} = the constrained optimal operating characteristic,

C_O = the optimal operating characteristic.

When it is necessary to refer to the i -th parameter case, a subscript will be added, giving d_{ci} .

Averaging d_{ci} emphasizes the constrained policy's performance by case and is used to measure its average response to changing parameters. This "casewise" average percent deviation is calculated as

$$(2.2) \quad \bar{d}_C = \frac{\sum_{i=1}^n d_{ci}}{n} = (1/n) \sum_{i=1}^n \left(\frac{C_{CO,i} - C_{O,i}}{C_{O,i}} \right) * 100$$

where n is the relevant number of observations.

An alternative method, however, aggregates the operating characteristic across the specified parameter set (for both the optimal and approximation policies), and then computes the average percent deviation

$$(2.3) \quad D_C = \frac{\sum_{i=1}^n C_{CO,i} - \sum_{i=1}^n C_{O,i}}{\sum_{i=1}^n C_{O,i}} * 100.$$

This method views the parameter set as an inventory system, each case being a different item stocked by the system. This "system perspective" emphasizes how the constrained optimal policy performs when applied over a large number of items, and is commonly used by inventory practitioners.

To clarify the difference between these two methods, consider a set of optimal and constrained optimal policies both resulting from the same parameter set. Computing the average percent deviation on a "casewise" basis \bar{d}_C prevents a poor performance for some cases with

small expected costs from being swamped by other cases with a small difference for larger expected costs. The inventory manager, however, wants to measure an inventory system's performance with each item contributing in proportion to its costs D_C . Thus, the average percent deviation is computed in two different ways depending upon what aspect of performance is being evaluated.

2.3 Initial Results

A full factorial set of the parameter values in Table 2.1 comprises 480 different cases, and an optimal (s, S) policy is computed for each case. Table 2.2 describes the distribution of s , S , and D , showing their percentile values. The reorder point s ranges from -4 to 10, and half of the time $s \geq 0$. The stockage objective S ranges from 0 to 24, and half of the time $S \geq 6$. (Note - one policy was $(-1, 0)$, meaning the item represented by that parameter case would be ordered as required and never stocked on-the-shelf). For $D = S - s$, the range is from 1 to 22, and half of its values are greater than or equal to 6. A list of all 480 policies is given in Appendix I. In 38% of the cases (182 out of 480), an optimal reorder point is negative. Also, the distributions of S and D values are similar, reflecting the generally small reorder point values. Let $C_{O,i}$ represent the expected total cost per period of an optimal policy for case i (where $i = 1, \dots, n$), and define

$$\bar{C}_O = \sum_{i=1}^n C_{O,i}/n. \text{ For } n = 480 \text{ cases, } \bar{C}_O \text{ is } \$2.06 \text{ and the average stockout frequency per item is } 0.10.$$

An important influence on \bar{C}_O is the service level value (80% or 90%). The higher the service level, the higher the expected cost per period. For the 240 cases with an 80% service level, \bar{C}_O is \$1.92 and the average stockout frequency per item is 0.14. For the 240 cases with

Table 2.2
Optimal (s,S) Policy Parameters - 480 Cases

Policy Parameter	Percentile				
	0	25	50	75	100
s	-4	-1	0	1	6
S	0	4	6	9	24
D	1	4	6	9	22

a 90% service level, C_o is \$2.20 and the average stockout frequency is 0.06. Most of the additional costs for the 90% service level come from holding more inventory. In other words, a 10% increase in the service level means increasing inventory investment by 33%.

The stock replenishment frequencies for the items ranged from 1 to 15 orders per year, with an average of 5 orders per year. The inventory turnovers per year ($\mu * 52 / \text{average on hand inventory}$) ranged from 2 to 42 times, and averaged 6 turnovers per year. The larger inventory values generally occurred when s was zero or negative. With the range of mean demand studied, these s values allowed the on hand inventory level to be zero for many periods prior to reordering, thus accruing no holding costs.

2.4 The Impact of Nonnegative Reorder Points

Note that 38% of the 480 optimal reorder points are negative. In one case, for example, the (s,S) policy was (-2,5). Thus, after reordering, the stockage level is five units. This gradually diminishes in future periods. A stock replenishment order is not placed until the first period after two backorders are recorded. Since $\mu = 0.1$ in this

case, on average, it will take 20 periods of "empty shelves" and "aging backorders" before a replenishment order is submitted!

Although nearly four times out of ten in the 480 item system it is optimal to wait until backorders are "on-the-books" before reordering, few firms can afford to operate in this manner. This is especially significant in a U. S. Air Force context, where the retail stockage model precludes negative reorder points [44].

2.4.1 The Constrained Policy

To evaluate the impact of ruling out negative reorder points, the (s, S) periodic review inventory model was modified to allow only nonnegative reorder points. The Veinott-Wagner algorithm was modified accordingly in Kaufman's program to compute the best (s, S) policy having a nonnegative reorder point. This was done by limiting the search range for S to values greater than or equal to D ; thus, $s = (S - D) \geq 0$. The algorithm, in computing the best "constrained" policy for a given parameter set, evaluated the 182 cases in which a negative reorder point was previously optimal.

2.4.2 Interpretation of the results

Table 2.3 compares the (s, S) policy parameters for the optimal (o) and constrained optimal (co) policies. For every percentile, the constrained s and S values are at least as large or larger than the optimal values. For the D values, however, this pattern is reversed. In either case, the difference between the optimal and constrained optimal S and D distributions is slight. Because the reorder points differ, however, optimal policies do have larger D values since their reorder points are negative.

Table 2.3

A Comparison of Optimal and Constrained Optimal (s,S)
Policy Parameters - 182 Cases

Policy Parameter	Percentile				
	0	25	50	75	100
s_o	-4	-1	-1	-1	-1
s_{co}	0	0	0	0	0
S_o	0	3	5	9	20
S_{co}	1	3	6	9	21
D_o	1	4	6.5	10	22
D_{co}	1	3	6	9	21

The 182 item system operating characteristics reveal a more pronounced difference between the two classes of policies (Table 2.4). Constrained policies allow fewer stockouts than optimal policies (0.03 versus 0.11). Essentially, constrained policies reorder more frequently, and hence order slightly fewer units than optimal policies, keeping more units on hand and accruing additional holding costs. Thus, on average, constraining an item's reorder point increases its total costs by 7.7%.

Table 2.4

A Comparison of Optimal and Constrained Optimal Operating Characteristics - 182 Item System

Policy	Average				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Optimal	\$1.68	\$.75	\$.68	\$.24	0.11
Constrained Optimal	1.81	.85	.89	.06	0.03

Table 2.5 provides more detail on the impact of parameter settings and the average percent changes due to the constrained policy. With increasing setup cost, constrained policies become relatively less costly, allowing only slightly more stockouts. As holding cost increases, constrained policies become relatively more costly, allowing

Table 2.5

Parametric Analysis of the Deviation Between Optimal and Constrained Optimal Policies When the Optimal Reorder Point is Negative - 182 Cases

Policy	Average Percent Deviation \bar{d}_c				No. Cases	
	Costs			Stockout Frequency		
	Total	Setup	Hold			
182 Items	10	18	41	-78	-76	182
K 3	14	32	62	-78	-77	56
20	8	11	32	-77	-75	126
h 0.1	8	11	30	-78	-75	67
0.3	10	15	47	-77	-75	47
0.5	11	23	48	-78	-76	38
0.7	13	28	50	-79	-77	30
p 4*h	11	17	46	-77	-73	122
9*h	7	19	30	-81	-79	60
L 0	12	19	43	-83	-81	123
2	6	16	40	-71	-67	43
4	4	11	34	-62	-57	16
μ 0.1	21	32	73	-87	-87	35
0.2	13	18	56	-83	-82	25
0.3	9	15	45	-79	-77	22
0.4	8	12	34	-76	-74	19
0.5	6	17	24	-73	-70	18
0.6	5	12	27	-74	-71	16
0.7	4	13	23	-72	-68	15
0.8	4	10	24	-69	-65	13
0.9	5	13	19	-74	-69	10
1.0	5	9	24	-73	-68	9

even fewer stockouts than optimal policies. With increasing penalty cost, however, constrained policies become relatively less costly while still allowing fewer stockouts than optimal policies. As the lead time lengthens, total cost and stockout frequency differences between the two classes of policies diminish. Similarly, as the mean demand grows, the total cost difference between the two classes of policies shrinks. Finally, negative reorder points are closely associated with low parameter settings (except for setup cost).

2.5 The Impact of Misspecifying Mean Demand

Here we begin to study how sensitive (s, S) policies are to errors in estimating (misspecifying) the mean demand per period μ . For practitioners, this is an important issue since inventory systems rely on statistical estimates of demand parameters. Often, μ is estimated as the sample mean \bar{x} (the number of observed demands n divided by the number of time periods t). If μ is misestimated, then the inventory policy used may not be correct and the firm pays a cost penalty for using an erroneous policy due to a misspecified mean. The probability of observing n units of demand over t periods depends on μ , and is calculable from the Poisson distribution with a mean $m = t\mu$. In this section, the cost penalties from over- or underestimating mean demand are related to the likelihood of that misspecification occurring.

2.5.1 Preliminary Investigation

By way of example, consider the case $\omega = 5/26$, $L = 4$, $h = 0.1$, $p = 4h$, and $K = 20$. Let $s(\cdot)$ and $S(\cdot)$ be constrained optimal values of (s, S) given mean demand μ and these other parameter values. Also let $A(s, S)$ be the associated expected cost when the (s, S) policy is used. For $\mu = 5/26$, $s(5/26) = 0$ and $S(5/26) = 9$, and $A(0, 9) = \$0.87$. The set

of misspecified means is $X = \{1/26, \dots, 15/26\}$. The corresponding constrained optimal (s, S) policies were calculated; for example, $s(12/26) = 0$ and $S(12/26) = 14$, and $A(0, 14) = \$0.96$. Consequently, the deviation in total cost from misspecifying the mean as 12/26 rather than 5/26 is $(0.96 - 0.87)/0.87 * 100 = 10\%$. Similarly, the other deviations for values in X were calculated and the results are plotted in the graph on the left of Figure 2.1. The horizontal axis represents the values $e = x - 5/26$, where $x \in X$. Define $F(e)$ as this curve.

Next we calculate the probability that $F(e)$ does not exceed a value V . Define $P(j) = \text{Poisson}(j | 5*26/26)$ for $j = 1, \dots, 15$ and $t = 26$, then we can calculate the desired probability as

$$(2.4) \quad \text{Prob}(F(e) \leq V) = \sum_{j \in W} P(j)$$

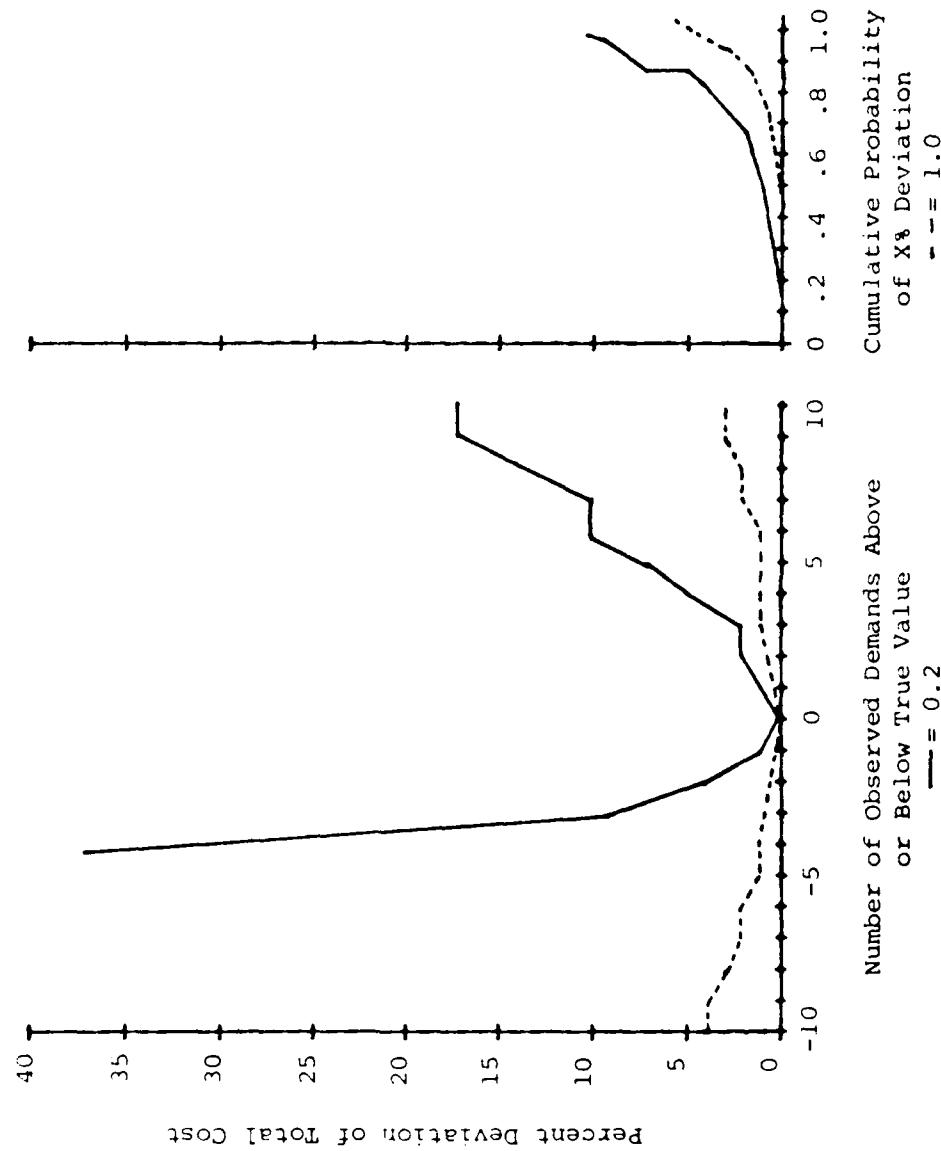
where $W = \{j | F(j/26) - 5/26 \leq V\}$. This probability function is then graphed in the right side of Figure 2.1.

2.5.2 Interpretation of the Results

Misspecifying the mean demand value increases the (s, S) policy's total system costs. Figure 2.1 compares the percent deviation in total costs for $\mu = 0.2$ and $\mu = 1.0$ demands per period. Consider the case where the true mean is 0.2 demands per period, roughly five demands every 26 periods. Suppose, however, that the mean is overestimated, that is, $\bar{x} = 15/26$ due to 15 demands being observed in 26 periods. Figure 2.1 shows that using a mean demand value almost three times larger than μ increases total system cost nearly 17% over optimum.

In general, Figure 2.1 shows the percent deviation in total costs from the optimum value is much larger for the lower mean demand value.

Figure 2.1
The Deviation of System Costs



Further, underestimating the true mean demand is worse than overestimating it.

The actual cost of misspecification becomes larger as the true mean increases. The cost of misspecification as a percent of the optimal cost, however, diminishes as the true mean increases. In essence, since the total cost for a low mean demand is small, there is greater room for error that significantly increases the percent deviation in costs.

Figure 2.1 also demonstrates this point probabilistically.

Extending a line from any given percent deviation in total cost to the desired cumulative probability curve reveals the fraction of time that percent deviation is not exceeded. For example, if the true mean is 0.2 demands per period, there is an 0.82 probability that the deviation in total costs will be 5% or less. Similarly, if the true mean is 1.0 demand per period, there is a 0.99 probability that the deviation in total cost will be 5% or less. Thus, for low mean demand items the impact of misspecifying mean demand is of moderate or little importance.

2.6 Conclusions

This chapter establishes three important findings for the parameter set examined. First, under conditions of low mean demand, an optimal (s, S) policy has a high likelihood of having a negative reorder point; for the parameter set in this study, 38% of the cases have negative s . Next, for the cases in this study with negative optimal s , constraining the (s, S) policies to nonnegative reorder points increases their average total system costs 8% over the optimum, mainly due to ordering more frequently and holding more inventory. The constrained policies have, on the average, 73% less stockouts than optimum policies.

Finally, a preliminary investigation suggests that constrained (s, S) policies may not be overly sensitive to misspecifying the mean demand. Even with means as small as 0.2 demands per period - when (s, S) policies are most sensitive to misspecification - there is an 0.82 probability that the additional total costs are within 5% of the optimal costs.

The findings in this chapter show that a naive application of (s, S) inventory policies leads to unexpected results for the parameter set studied. Frequently, these results suggest employing negative reorder points, essentially waiting until backorders are "on-the-books" before reordering stock. This logic, however, is not acceptable to many organizations, including military supply centers at logically remote locations. The problem of negative reorder points may be resolved by constraining s to nonnegative values, incurring only modest additional inventory costs. Further, a preliminary look at the cost penalties arising from misspecifying demand parameters suggests that, if demand is Poisson distributed, the penalties are not excessive. With these encouraging results, Chapter 3 searches for a method of finding constrained (s, S) policies that eases the computational burden of computing optimal policies. Chapter 4 then takes this approximation technique and evaluates its performance when demand parameters are statistically-derived.

CHAPTER 3

A COMPARISON OF CONSTRAINED APPROXIMATELY OPTIMAL (s, S) INVENTORY POLICIES UNDER LOW DEMAND (KNOWN MEAN)

3.1 Introduction

As was shown in Chapter 2, negative reorder points are common under low demand, and additional costs are incurred when operating with constrained (s, S) policies. Although constrained optimal policies are not overly sensitive to misspecified mean demand values, they are burdensome to calculate.

Approximation methods offer attractive computational alternatives to finding optimal policies. Approximation methods provide near-optimal inventory system performance and are less complex computationally. Two well-known approximation methods are Ehrhardt and Mosier's Revised Power Approximation [15] and Naddor's "Analogy" Approximation [34]. The performance of these approximations, however, has not been examined for the low mean demand and constrained reorder point environment.

This chapter evaluates the performance of these two approximations operating with constrained reorder points and infrequent demand. In this chapter, the value of mean demand is assumed to be known; the assumption is dropped in Chapter 4. We demonstrate that the Analogy Approximation and Power Approximation perform equally well for the parameter set in this study. For the 480 item system described in Chapter 2, the constrained optimal policies' average total cost per

item is \$2.11 as compared with \$2.13 for both constrained approximations. On average, however, the Analogy Approximation policies' total cost deviates 0.7% from the constrained optimum whereas the constrained Power Approximation policies' total cost deviates 1.1%. This near-optimal performance of the Power Approximation is reaffirmed with parameter values that are interpolated and extrapolated from the original parameter set. Comparing constrained optimal and Power Approximation policies for these 32 new parameter values reveals that the average total cost performance per item (\$1.84 versus \$1.92) and the average deviation in total cost (5.3%) are somewhat worse than with the 480 item system. Nevertheless, this performance is acceptable from a practitioner's viewpoint.

Section 3.2 describes the parameter set, and how to compute both approximately optimal policies. Section 3.3 evaluates the two approximation methods in the 480 item system, Section 3.4 examines the Power Approximation for the interpolated and extrapolated parameter values, and Section 3.5 draws conclusions from these results.

3.2 Methodology

3.2.1 The Power Approximation

The Revised Power Approximation of Ehrhardt and Mosier [15] "is an algorithm for computing approximately optimal values for (s, S) using only the mean and variance of demand [15]."
The algorithm's computations proceed as follows.

Recall that K is the ordering setup cost, h is the holding cost per unit per period, p is the penalty cost per unit per period, L is the delivery lag, and μ is the mean demand per period. Then calculate
(3.1) $D_p = (1.3)\mu^{.494}(K/h)^{.506}(1+(L+1)\sigma^2)/\mu^2^{.116},$

and

$$(3.2) \quad s_{p1} = .973(L+1)\mu + (L+1)^{.5}\sigma(.183/z + 1.063 - 2.192z),$$

where z is found by

$$(3.3) \quad z = D_p / ((p/h)(L+1)^{.5}\sigma)]^{.5}.$$

Finally, calculate

$$(3.4) \quad s_p = \max(s_{p1}, 0),$$

and

$$(3.5) \quad S_p = s_{p1} + D_p.$$

The constrained Power Approximation policy is (s_p, S_p) . Note that (3.4) ensures $s_p > 0$.

Equations (3.1) to (3.5) are appropriate whenever $D_p/\mu > 1.5$, which is true for the parameter sets examined in this study.

When $D_p/\mu \leq 1.5$ then calculate

$$S_2 = (L+1)\mu + v[(L+1)\sigma^2]^{.5},$$

where v is the solution to

$$\phi_N(v) = p/(p+h),$$

and $\phi_N(v)$ is the cumulative unit normal distribution function.

Then $s_p = \min(s_{p1}, S_2, 0)$ and $S_p = \min(S_p, S_2)$.

The values in (3.4), (3.5), and S_2 are rounded to the nearest integer since demand is assumed to be integer valued.

Observe that D_p could be rounded instead of s_p [15]. But rounding D_p increased the approximation's average percent deviation from optimal costs (typically from 0.2% to 0.4%), and so S_p was rounded.

3.2.2 The Analogy Approximation

The Analogy Approximation computes (s, S) policies by taking s and S values from "analogous" policies [34,35]. The computations are similar to those described in Porteus [35]. Specifically,

$$(3.6) \quad n = \max (1, (2K(p+h)/(ph))^{.5}),$$

and

$$(3.7) \quad Q = (2K\mu(p+h)/(ph))^{.5}.$$

The reorder point s_a is found by

$$(3.8) \quad s_a = \max((L+1)\mu + (1-Q)/2 + \alpha(\sigma^2(L+1)+(Q^2-1)/12)^{.5}, 0),$$

where α is

$$(3.9) \quad \alpha = \Phi_N^{-1}(p/(p+h)).$$

The symbol Φ_N^{-1} denotes the inverse of the cumulative unit normal distribution. Let

$$(3.10) \quad \beta = L + (n+1)/2,$$

and calculate

$$(3.11) \quad S_a = \beta\mu + \alpha(\beta\sigma^2 + (n^2-1)\mu^2/12)^{.5}.$$

Now, $S_a = \max(s_a, S_a)$. Furthermore, the values for n , Q , s_a , and S_a are rounded to the nearest integer as they are calculated.

Note that in (3.11) the variance of demand per period is used, rather than the variance of demand over lead time as reported in Porteus [21]. This modification agrees with Ehrhardt and Kastner's derivation of a method to calculate the Analogy Approximation [12]. (This correction to Porteus is also pointed out by Jacobs [26]).

3.2.3 Parameter Set

The first half of the research experiment examines the performance of two approximation methods and uses the 480 item system described in Section 2.2.4. The second half of the experiment evaluates the sensitivity of the Power Approximation method to a different set of parameter values. Each of the five parameters (K , h , p , L , and μ) is given two new values, one between the original values (interpolated) and the other outside the range of the original values (extrapolated).

This results in the full factorial design with 32 cases shown in Table 3.1.

Table 3.1
Experiment with 32 Interpolated and Extrapolated
Parameter Cases

Parameter	Values
K	\$5, 35
h	.4, 1.0
p	5*h and 12*h
L	1, 6
μ	0.05, 0.15 (Poisson)

3.2.4 Evaluating Performance

The performance measures are the average percent deviations from the optimum for the differing (s,S) policy operating characteristics. The averages, however, are computed in two different ways, depending upon what aspect of performance is being evaluated. Recall from Section 2.2.6 that the casewise average percent deviation emphasizes individual differences, whereas the aggregate average percent deviation emphasizes inventory system performance.

3.3 Evaluation of Results

3.3.1 Approximation Results

Four sets of approximately optimal policies, the Power and Analogy Approximations for both constrained and unconstrained reorder points, are calculated for each of the 480 cases in Table 2.1. Kaufman's

optimization algorithm [28] was used to determine the operating characteristics for each case.

Table 3.2 compares the (s, S) policy parameters for the 480 constrained optimal (co), constrained Power Approximation (cp), and constrained Analogy Approximation (ca) cases. Overall, the distribution of policy parameters s , S , and D for the three different policies are similar; however, there are three points of difference. First, the Power Approximation has the largest range of S and D values (25=26-1 and 23=24-1, respectively). Second, both approximations' S values are slightly larger than optimal in the 25th through 75th percentiles. Finally, the Analogy Approximation's D values are slightly larger in

Table 3.2

A Comparison of Constrained Optimal with the Constrained Power and Analogy Approximations' Policy Parameters - 480 Cases

Policy Parameter	Percentile				
	0	25	50	75	100
s_{co}	0	0	0	1	6
s_{cp}	0	0	0	2	6
s_{ca}	0	0	0	1	6
S_{co}	1	4	6	9	24
S_{cp}	1	4	7	10	26
S_{ca}	1	4	7	10	24
D_{co}	1	3	5	8	22
D_{cp}	1	3	5.5	8	24
D_{ca}	1	3	6	9	22

Legend:

s = Reorder Point
 S = Stockage Objective
 D = $S-s$

co = Constrained Optimal
 cp = Constrained Power Approx.
 ca = Constrained Analogy Approx.

the 50th through 75th percentiles than for the optimal or Power Approximation policies. The consequences of these differences are shown in Table 3.3.

Table 3.3

A Comparison of Constrained Optimal with the Constrained Power and Analogy Approximations' Operating Characteristics -
480 Item System

Policy	Average				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Optimal	\$2.11	\$.81	\$1.03	\$.27	0.07
Power Ap.	2.13	.81	1.11	.21	0.06
Analogy Ap.	2.13	.78	1.12	.23	0.06

On average, the total cost per item is \$2.13 for both approximations, compared to \$2.11 for the optimum. The stockout frequency for both approximations (0.06) is lower than the optimum (0.07).

Next, we compare the Power and Analogy Approximations. The Power Approximation orders slightly smaller amounts more frequently than the Analogy Approximation. Thus, it experiences stockouts of a shorter duration than the Analogy Approximation. Also, the Analogy Approximation holds slightly more inventory than the Power Approximation.

The differences between the two approximations are clearer in Table 3.4, showing the average percent deviation for a 480 item system. The total cost per item for the Analogy and Power Approximations average 0.7% and 1.1% higher than optimal, respectively. This compares to a

0.5% deviation from optimum obtained by Ehrhardt and Mosier [15] for the unconstrained Revised Power Approximation.

Table 3.4

**Average Percent Change in Operating Characteristics
Over the Optimum for Constrained Power and Analogy
Approximation Policies Given a 480 Item System**

Policy	Average Percent Change				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Power Ap.	1.1	0.1	8.3	-23.2	-17.1
Analogy Ap.	0.7	-3.8	9.0	-17.1	-14.2

The Power Approximation also reduces stockouts 17% below optimum; the Analogy Approximation reduces stockouts only 14%. In essence, the Power Approximation trades off higher setup costs for lower holding and penalty costs whereas the Analogy Approximation does the opposite.

Further insights into the approximations' performance over the 480 cases cases are derived by grouping the cases as follows:

Block A: $s_o > 0$ and $s_p > 0$ (or $s_a \geq 0$),

Block B: $s_o > 0$ and $s_p < 0$ (or $s_a < 0$),

Block C: $s_o < 0$ and $s_p > 0$ (or $s_a \geq 0$),

Block D: $s_o < 0$ and $s_p < 0$ (or $s_a < 0$),

where the subscript o designates optimal policies, the subscript p designates unconstrained Power Approximation policies, and the subscript a designates unconstrained Analogy Approximation policies.

When $s \geq 0$, the associated unconstrained policy's operating characteristics were used; whenever $s < 0$, the associated constrained policy's operating characteristics were used. Since the Power and

Analogy Approximations may recommend reorder points in opposite directions of optimal s for the same parameter case, the number of observations for each approximation in the blocks need not be identical.

Table 3.5 gives the average percent deviation \bar{d}_c for each of the four blocks' operating characteristics. For example, 297 Power Approximation cases are in Block A, and their casewise average percent deviation from optimal total costs is 1.8%. Similarly, 157 Analogy

Table 3.5

Casewise Average Percent Difference Between Operating Characteristics for Constrained Power and Analogy Approximation Policies and the Optimum - 480 Cases

Policy (Cases)	Casewise Average Percent Change				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Block A					
Power (297)	1.8	3.9	14.1	-22.3	-21.3
Analogy (293)	1.0	- 2.8	13.4	-15.6	-14.9
Block B					
Power (1)	0.1	- 4.5	6.7	- 4.5	- 4.5
Analogy (5)	0.0	- 1.2	1.4	- 1.2	- 1.2
Block C					
Power (80)	0.4	- 6.8	8.2	- 6.7	- 6.7
Analogy (25)	4.2	-23.7	30.8	-23.7	-23.7
Block D					
Power (102)	0.5	- 0.9	2.4	- 0.9	- 0.9
Analogy (157)	0.3	- 4.6	5.6	- 4.6	- 4.6

Definition of Blocks:

- Block A: $s_o \geq 0$ and $s_p \geq 0$ (or $s_a \geq 0$),
- Block B: $s_o \geq 0$ and $s_p < 0$ (or $s_a < 0$),
- Block C: $s_o < 0$ and $s_p \geq 0$ (or $s_a \geq 0$),
- Block D: $s_o < 0$ and $s_p < 0$ (or $s_a < 0$).

Approximation cases are in Block D, and their average percent deviation from optimal total costs is 0.3%.

From Table 3.5 we see there are 399 Power Approximation policies and 450 Analogy Approximation policies in Blocks A and D. Thus, the Analogy Approximation recommends a reorder point with the same sign (+ or -) as optimal more frequently than the Power Approximation. Further, there are only five Analogy and one Power Approximation policies that recommend a negative reorder point when the optimal policy recommends a positive one (Block B). Finally, in Blocks A, B, and D the Analogy Approximation's average percent deviation in total cost per item is 1.0%, 0.0%, and 0.3% (respectively), compared to the Power Approximation's 1.8%, 0.1%, and 0.5%. For these blocks, the Power Approximation's higher average percent deviation is partially explained by its relatively more frequent ordering (especially for Block A).

3.3.2 Interpretation of the Results

For the 480 item system, the Power and Analogy Approximations have nearly identical average total costs (\$2.13) and stockout frequencies (0.06). On an aggregate basis, however, the Analogy Approximation deviates less from optimal total costs than the Power Approximation (0.7% versus 1.1%). Further, although both approximations allow fewer stockouts than optimal, the Power Approximation provides greater protection from stockouts than the Analogy Approximation (-17.1% versus -14.2%).

These results are explained by the tradeoffs each approximation is making. The Analogy Approximation holds 9% more inventory yet orders 3.8% less frequently than optimal. In contrast, the Power Approximation holds 8.3% more inventory yet orders 0.1% more frequently than optimal.

Thus, although the Analogy Approximation deviates less from optimal total costs, the Power Approximation provides better service. The most remarkable result, however, is that both approximations performed close to optimal over a large parameter set.

3.4 Interpolation and Extrapolation Analysis

The Power Approximation was used to test the sensitivity of approximately optimal policies to parameters interpolated and extrapolated from the original parameter set. Table 3.6 shows the distribution of policy parameters for the constrained optimal and Power Approximation policies. From Table 3.6, we see that the ranges for s_{CO} ($2=2-0$), s_{CP} ($7=8-1$), and D_{CO} ($6=7-1$) are larger than the corresponding ranges for s_{SO} ($1=1-0$), s_{SP} ($5=6-1$), and D_{SP} ($5=6-1$). Further, 50% of the s_{SP} and D_{SP} values are between 2 and 4, whereas 50% of the s_{CO} and D_{CO} values are between 1.5 and 3.

Table 3.6

A Comparison of Constrained Optimal and Constrained Power Approximation Policy Parameters - 32 Interpolated and Extrapolated Cases

Policy Parameter	Percentile				
	0	25	50	75	100
s_{CO}	0	0	0	0	1
s_{SP}	0	0	0	0.5	2
s_{SO}	1	1.5	2.5	3	6
s_{SP}	1	2	3	4	8
D_{CO}	1	1.5	2	3	6
D_{SP}	1	2	3	4	7

The pattern of approximately optimal policies exceeding optimal cost found in Section 3.3.1 applies for the 32 item system as well. On average, the Power Approximation is more expensive (\$1.92 versus \$1.84) yet allows fewer stockouts (0.02 versus 0.03) than optimal (Table 3.7). The higher cost of approximation policies is due to a nearly 30% higher average inventory investment than optimal (\$1.30 versus \$1.00). Even the resulting savings in ordering and stockout penalty costs cannot offset the additional costs from the approximation's higher inventory investment.

Table 3.7

A Comparison of Constrained Optimal and Constrained Power Approximation Operating Characteristics - 32 Item System

Policy	Average				Stockout Frequency	
	Costs			Penalty		
	Total	Setup	Hold			
Optimal	\$1.84	\$.64	\$1.00	\$.20	0.03	
Power Ap.	1.92	.51	1.30	.11	0.02	

The most interesting result, however, is found in Table 3.8, which summarizes the approximation's casewise percent deviation from optimal total costs. Two-thirds of the 32 cases deviate from the optimal by less than 5%. Further, seven of the eleven cases with 5.0% or more deviation had a mean demand value of 0.05 demands/period (roughly 3 demands/year). In fact, the four cases where the approximation performed the worst had a mean demand of 0.05 demands/period. Nonetheless, viewing the cases as a 32 item system, the approximation's average percent deviation from optimal total costs is only 4.6%. This

indicates how robust approximately optimal policies are to interpolated and extrapolated parameters.

Table 3.8

The Deviation From Optimal Total Costs for the Power Approximation as Applied to 32 Interpolated and Extrapolated Parameter Cases

Percent Deviation in Total Costs	Number of Cases	Cumulative Percentage
[0.0, 1.0)	8	25%
[1.0, 3.0)	6	44%
[3.0, 5.0)	7	66%
[5.0, 10.0)	6	84%
[10.0, 15.0)	2	91%
[15.0, 20.0)	1	94%
[20.0, 25.0)	1	97%
[25.0, 26.0]	1	100%

3.5 Conclusions

This research has compared the performance of constrained Power and Analogy Approximations to constrained optimal (s, S) policies. The central result is how close to optimal both approximations performed for the 480 test cases. On average, for a 480 item system, both approximations deviated from the optimal total costs by roughly 1% or less. Recall that this compares with the 0.5% deviation reported by Ehrhardt and Mosier [15] for the unconstrained Revised Power Approximation. In essence, the Power and Analogy Approximations' higher total costs come from a different balance being struck among the

components of total cost. That is, the constrained approximations trade slightly higher holding costs for better service.

The approximation performance deteriorates slightly with the interpolated and extrapolated parameter set. On average, total costs are 4.6% higher than optimal. Many of the cases with total cost deviations exceeding 5% are for a mean demand of less than three demands per year (the smallest value used in this research). These results are still acceptable from a practitioner's point of view.

The findings in Chapters 2 and 3 make two significant points for inventory managers. First, Chapter 2 shows that negative reorder points in (s, S) inventory policies may be eliminated by constraining $s > 0$. Although constrained reorder points increase the system's average total costs, the effect is not excessive. For the 182 items with constrained reorder points, the increase averaged only 7.7%. Chapter 3 has shown that constrained (s, S) policies are effectively computed using approximation techniques. The cost penalty for using approximately optimal (s, S) policies is small. The results in this chapter assume that the mean demand per period is known. Chapter 4 relaxes this assumption and reevaluates the performance of approximately optimal policies when demand parameters are statistically-derived.

CHAPTER 4

AN EVALUATION OF CONSTRAINED APPROXIMATELY OPTIMAL (s,S) INVENTORY POLICIES UNDER LOW DEMAND (ESTIMATED MEAN)

4.1 Introduction

Approximately optimal (s,S) policies diminish the computational burden of calculating optimal (s,S) policies. Chapter 3 showed that approximation methods perform well under low demand where the mean is known. For an experimental system of 480 items, the overall increase in cost from using approximately optimal policies ranged from 0.7% to 1.1%

A critical test of the approximation methods is how well they perform in a statistical demand environment, where the demand parameters are estimated from historical demand data. This test is important since even "optimal" (s,S) policies may not be optimal when their demand parameters are estimated. As discussed in Chapter 1, finding an optimal statistical estimation procedure remains an open research issue. Thus, a computer simulation model is used to evaluate the approximations in a statistical environment.

This chapter assesses the performance of constrained Power and Analogy Approximations when mean demand per period is low and is statistically estimated from historical data. The research establishes that both approximations perform similarly. For a 96 item system, the constrained optimal policies' average total cost per item is \$1.95 as compared to \$2.01 and \$2.02 for the constrained Analogy and Power Approximations, respectively.

The differences between the two approximations are more distinct comparing their aggregate deviation from optimal policy operating characteristics (equation 2.3). On average, the Analogy Approximation's total cost is 3.1% greater than optimal while the Power Approximation's total cost is 3.8% larger. This difference is statistically significant. As in Chapter 3, both approximations have fewer stockouts than optimal; the Power Approximation provides better service than the Analogy Approximation (a reduction in stockouts of -13.4% versus -9.7%), but requires a higher inventory investment (an increase of 11.1% versus 10.0%).

Section 4.2 reviews the approximation methods, the simulation model's structure, the test parameter set, and the performance criteria. Section 4.3 evaluates the simulation results for two 96 item systems operating under each approximation, one has Poisson demands of the other Negative Binomial demands. Finally, Section 4.4 draws conclusions from these results.

4.2 Simulation Method

4.2.1 The Approximation Techniques

The constrained Power and Analogy Approximations, described in Sections 3.2.1 and 3.2.2, respectively, are used in the simulations.

4.2.2 Description of the Simulation

The simulation is accomplished by a FORTRAN program that evaluates approximately optimal (s,S) inventory policies for multi-item systems. It was written by MacCormick in 1975 and refined by Ehrhardt in 1976. The simulation program computes the expected operating characteristics for each item in the system [10,11], using conditional Monte Carlo sampling and antithetic sampling techniques to reduce the variance of

these estimates [14]. Many other researchers have used this simulation, including MacCormick [32], Estey and Kaufman [17], Klincewicz [30], Ehrhardt [10,11], Kaufman [29], Blazer [4], and Jacobs [26].

Three important modifications were made to the simulation model for this research. First, since demand is Poisson distributed, the subroutine estimating the demand parameters was modified to set the sample variance equal the sample mean. Second, the subroutines calculating the inventory policies were modified to compute the constrained Power and Analogy Approximations. Finally, the simulation was modified to accommodate mean demand values smaller than one demand per period. These modifications were validated by: 1) printing intermediate values, 2) computing policies with the same source code used in Chapter 3 and verified with hand calculations, and 3) comparing the simulation output with optimal policies and their operating characteristics for reasonableness.

4.2.3 Experimental Design

Due to the simulation's lengthy run time, a subset of 96 cases, shown in Table 4.1, was selected from the 480 case parameter set used

Table 4.1
Experiment with 96 Parameter Cases

Parameter	Values
K/h	\$4.3, 10, 40, 200
h/h	\$1
p/h	\$4, 9
L	0, 2, 4
μ	0.2 (.2) 0.8 (Poisson)

in Chapters 2 and 3. These 96 cases are broadly representative of the original parameter set. Note that K , p and h are reexpressed as K/h , p/h and $h/h=1$ in Table 4.1. Normalizing the value of h is required by the simulation's design.

For each parameter case, the approximately optimal (s, S) policy is recomputed every 26 periods, with the simulation running for 200 such revisions. Thus, if the period is a week, the (s, S) policy is revised semiannually and the simulation runs for 100 years. At each policy revision, the previous 26 periods of demand history are used to compute a new sample mean for demand. The demand estimate is then used to recompute the (s, S) policy to be used for the next 26 periods.

Further, from the 100 years of simulated history, the model computes several estimated operating characteristics, including: 1) expected total cost per period, 2) expected setup cost per period, 3) expected holding cost per period, 4) expected penalty cost per period, and 5) stockout frequency. These operating characteristics are directly comparable to outputs for optimal policies produced by Kaufman's program [28].

4.2.4 Evaluating Performance

The performance of constrained optimal (s, S) policies under known mean demand is the basis of comparison for the constrained Power and Analogy Approximations operating in a statistical environment.

Performance is measured as the average percent deviation from the optimum for the differing (s, S) policy operating characteristics, Section 2.2.6 described how this performance degradation is calculated. Note that the deviation may be calculated on a casewise or an aggregate base, depending on whether individual or system differences are to be emphasized.

4.3 Evaluation of Results

4.3.1 Simulation Results

Table 4.2 compares the distribution of (s, S) policy parameters for the 96 constrained optimal (co), constrained Power Approximation (cp) and constrained Analogy Approximation (ca) cases. For each of the 96 cases, the (s, S) policy parameters for both approximations were averaged over 100 years of simulated history; Table 4.2 shows the distribution of these average s , S , and D values. The constrained optimal (s, S) policy parameters are the s , S , and D values when mean demand is known. The distribution of policy parameters (s , S , and D) for both approximations in the simulation is similar to that for optimal policies. The most notable difference is in the range of S and

Table 4.2

A Comparison of (s, S) Policy Parameters for the Constrained Power and Analogy Approximations in a Statistical Environment with Constrained Optimal Policies with Known Mean Demand - 96 Cases

Policy Parameter	Percentile				
	0	25	50	75	100
s_{co}	0	0	0	1	5
s_{cp}	0	0	0	1.5	5
s_{ca}	0	0	0	1	5
S_{co}	1	4	6	9	21
S_{cp}	1	4	6	10.5	23
S_{ca}	2	4	6	10	21
D_{co}	1	3	5	9	20
D_{cp}	1	3	5	9.5	21
D_{ca}	2	3	5	9.5	19

Legend:

s = Reorder Point

co = Constrained Optimal

S = Stockage Objective

cp = Constrained Power Approx.

D = $S-s$

ca = Constrained Analogy Approx.

D values. The Power Approximation's S and D ranges (22=23-1 and 20=21-1, respectively) are larger than the optimal ranges (20=21-1 and 19=20-1), whereas the Analogy Approximation's ranges (19=21-2 and 17=19-2) are smaller than optimal. Section 3.3.1 notes that the Power Approximation's S and D ranges are larger than optimal under known mean demand too. In contrast, the Analogy Approximation, under known mean demand, had s, S, and D ranges identical to the optimal.

Table 4.3 shows the aggregate operating characteristics for the 96 item system. When the mean demand per period is known, the average total cost per item for the Analogy and Power Approximations is slightly greater than optimal (\$1.97 and \$1.98 versus \$1.95, respectively). Also, the approximations allow fewer stockouts than

Table 4.3

A Comparison of Constrained Optimal with the Constrained Power and Analogy Approximations' Operating Characteristics Under Known and Estimated Mean Demand - 96 Item System

Policy	Average				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Optimal (kn)	\$1.95	\$.72	\$.97	\$.27	0.07
Power (kn)	1.98	.72	1.06	.19	0.05
Power (est)	2.02	.74	1.07	.22	0.06
Analogy (kn)	1.97	.70	1.05	.22	0.06
Analogy (est)	2.01	.71	1.06	.24	0.06

optimal (0.06 and 0.05 versus 0.07). As in Section 3.3.1, the Power Approximation incurs higher ordering costs and lower stockout penalty costs than the Analogy Approximation. However, it also holds slightly more inventory than the Analogy Approximation.

The operating characteristics are aggregated for the statistically estimated mean demand per period (Table 4.3). The Analogy and Power Approximations have a higher average total cost per item than optimal (\$2.01 and \$2.02 versus \$1.95, respectively), yet allow fewer stockouts than optimal (0.06 and 0.06 versus 0.07). (Note that the optimal policy results assume that the mean demand is known). As before, the Power Approximation has higher ordering and holding costs than the Analogy Approximation, but lower stockout penalty costs.

The results in Table 4.4 show the aggregate average percent deviation in operating characteristics for the 96 item system. When the mean demand per period is known, the Analogy Approximation system averages 0.8% higher total costs than optimal, while the Power

Table 4.4

A Comparison of Operating Characteristics for Two Constrained Approximation Policies to the Optimum When Mean Demand is Known and Estimated - 96 Item System

Policy	Average Percent Deviation				
	Costs				Stockout Frequency
	Total	Setup	Hold	Penalty	
Power (kn)	1.3	0.8	9.7	-27.5	-18.9
Power (est)	3.8	2.3	11.1	-18.9	-13.4
Analogy (kn)	0.8	-2.9	8.8	-18.1	-13.6
Analogy (est)	3.1	-1.0	10.0	-10.5	- 9.7

Approximation system averages 1.3% higher costs. The Power Approximation allows 18.9% fewer stockouts than optimal, compared to the Analogy Approximation which allows 13.6% fewer stockouts.

When the mean demand per period is statistically estimated, both approximations realize additional performance degradation. The aggregate average percent deviation in total costs for both approximations is roughly three to four times larger than it was when mean demand is known. The Wilcoxon matched-pairs signed-rank test was used to determine if there was a statistically significant cost difference between the Analogy and Power Approximation policies. The Wilcoxon matched-pairs signed-rank test allows making inferences about treatments (Power and Analogy Approximations) when sample pairs are matched on all the other parameter values. The test is appropriate when the amount and direction of difference between the paired observations can be measured (total cost), and the observations are based on an element of randomness (simulated demand). The test revealed that there is a significant difference in the approximations' cost performance at a significance level of 0.01 ($z = -2.5$). The reduction in stockouts is noticeably less. Nevertheless, the increase in systemwide costs that results from using statistically estimated means is only 3 to 4% above optimal with known mean.

4.3.2 Supplementary Simulation with $\sigma^2/\mu = 9$

For a supplementary comparison of the constrained approximation methods, their relative performance is evaluated for a demand distribution having variance larger than the mean. The simulation is repeated using the parameter set in Table 4.1, except that demand is assumed to have a Negative Binomial distribution with a variance-to-mean ratio of nine. The selection of this variance-to-mean ratio for the Negative Binomial distribution is motivated by the earlier work of Ehrhardt [10,11]. In a low demand environment, the

Negative Binomial distribution may be appropriate when separate demand occurrences within a period are for multiple items. For example, replacement demands for aircraft tires at a base may evidence a high variance-to-mean ratio. For these supplementary simulations, the sample mean and sample variance of 26 historical demands are used to estimate the μ and σ^2 in the Power and Analogy Approximation methods at each policy revision.

Appendix II shows the 96 item system operating characteristics and their aggregate deviation from optimal values in a statistical environment. The Power and Analogy Approximations' average total cost per item is \$3.90 and \$4.02, respectively, as compared to the optimal value \$3.23. Thus, the Power Approximation incurs a 20.8% increase in cost as compared to 24.6% for the Analogy Approximation, when compared to optimal with known mean and variance. Using the Wilcoxon matched-pairs signed-rank test, this cost performance difference is statistically significant ($p < 0.001$, $z = -6.0$). The approximations also allow fewer stockouts than optimal. Thus, the increase in costs is primarily due to investing in higher levels of inventory.

4.3.3 Interpretation of Results

For the 96 item system, both approximations using statistically estimated mean demand experience higher total costs but fewer stockouts than optimal policies operating under known mean demand. The Power Approximation orders more frequently and holds more inventory than optimal. The Analogy Approximation orders less frequently, and holds more inventory than optimal. When demand is low, the Analogy Approximation has slightly smaller total costs than the Power

Approximation when demand is Poisson, but the Power Approximation is somewhat better when demand is Negative Binomial.

The most notable result is the aggregate average percent deviation from optimal total costs for the 96 item system with Poisson demand. It is shown in Table 4.4 that this aggregate average percent deviation is only 3.1% and 3.8% for the constrained Analogy and Power Approximations, respectively.

Ehrhardt [10] used the same inventory simulation with a 72 item system operating under the unconstrained Power Approximation. He compared Power Approximation policies under statistically estimated demand parameters with optimal (s, S) policies for the "full-information" case (known mean and variance of demand). Although his parameter set differed from Table 4.1 (including larger mean demand values from a Negative Binomial distribution), Ehrhardt reported an aggregate average deviation 11.5% above the optimal total cost [10]. In the supplementary simulation, given in Appendix II, the comparable percent rises to 20.8% when demand is low.

Jacobs [26] also used the same inventory simulation with intentionally biased, statistically estimated demand parameters. With virtually the same 72 item parameter set as Ehrhardt's, Jacobs compared the unconstrained Analogy and Power Approximations for the statistical and full-information environments. His biased estimation procedure reduced the degradation resulting from statistically-derived demand parameters. Jacobs reports aggregate average deviations 7.4% and 8.5% above the full-information Power and Analogy Approximations, respectively [26].

Thus, when it is reasonable to assume that demand is Poisson distributed, approximately optimal (s, S) policies perform extremely well in a statistical environment. This performance would be highly satisfactory to inventory practitioners.

4.4 Conclusions

This research has examined how the constrained Power and Analogy Approximations perform in a statistical environment. Overall, both approximations performed very well, deviating less than 4% from the constrained optimum while operating in a statistical demand environment. On average, both approximations are slightly more costly but allow fewer stockouts than optimal. Consistent with the previous chapter, the Analogy Approximation has a slight edge in total cost performance, and the Power Approximation provided better service. The difference in cost performance between the two approximations is statistically significant at the 0.01 level. When demand is Negative Binomial, however, the relative cost performance of the two approximations reverses. In this case the difference in cost performance is statistically significant at the 0.001 level.

In summary, it has been shown that approximately optimal (s, S) inventory models perform well when mean demand is low and the reorder point is constrained to be nonnegative. Furthermore, these constrained, approximately optimal inventory models perform to within 1% of optimal total costs when the mean demand is known, and 4% when the mean demand is statistically estimated from historical data. Inventory practitioners would find this degree of cost performance acceptable. The focus of this study now shifts to a search for approximating formulas for (s, S) inventory policy operating characteristics in Chapter 5.

CHAPTER 5

REGRESSION DERIVED APPROXIMATIONS FOR CONSTRAINED OPTIMAL AND APPROXIMATELY OPTIMAL (s,S) INVENTORY POLICY OPERATING CHARACTERISTICS

5.1 Introduction

Besides simplifying the calculation of (s,S) inventory policies, approximation techniques also lessen the associated information requirements. As with (s,S) approximations, decision makers can apply approximations to quickly and easily estimate the operating characteristics of inventory policies. This gives managers better insight into the likely operational impact of new inventory policies without having to actually implement them. These operational impacts include: gauging the sensitivity of average total cost to changing mean demand, or the impact of consolidating warehouses, or evaluating the benefit of shortening lead times.

This chapter develops regression derived operating characteristic approximations for the constrained optimal and constrained Power Approximation inventory policies under low demand. Since Chapter 3 showed both policies perform similarly, it is not too surprising that their operating characteristic approximations are also similar. The five approximations are: 1) total cost, 2) replenishment cost, 3) holding cost, 4) backlog cost, and 5) backlog protection. All the approximations, except for backlog cost, have coefficients of determination (r^2) of 0.944 or higher and average absolute deviations of less than 12% from their respective actual values. Further,

comparing the total cost approximations with previously reported results confirms that these approximations are dependent on the parameter set. That is, when approximations developed for a specific parameter set are used to predict the inventory policy operating characteristics of other parameter sets, the approximation's performance degrades if there is a substantial difference between the two sets.

Section 5.2 specifies the regression model, the associated parameter set, and the performance criteria that are fit. Section 5.3 describes the five performance approximations for the constrained optimal and Power Approximation inventory policies, and compares the two total cost approximations' performance with a previously reported approximation. Section 5.4 proposes several applications of the approximations. Finally, Section 5.5 draws conclusions about these operating characteristic approximations.

5.2 Methodology

5.2.1 The Regression Models

The operating characteristic approximation regression models are patterned after the general model described by Ehrhardt [16] in 1985. Following his notation, the operating characteristics which include total cost T , holding cost H , replenishment cost R , and backlog cost B , are transformed to T/h , H/h , R/h , and B/h , respectively. Further, backlog protection P is defined as one minus the stockout frequency. These characteristics are the dependent variables for a series of curvilinear, stepwise regression runs.

Recall that K is the ordering setup cost, h is the holding cost per unit per period, p is the penalty cost per unit per period, L is the

delivery lag, and μ is the mean demand per period. Given these parameters, the four cost characteristics (T/h , H/h , R/h , B/h) are fit to the same general multiplicative model which was shown to be effective in Ehrhardt [16]. For example, the expression for total cost is

$$(5.1) \quad T/h = c(L+1)^a(L+1) + b(1/L+1) + c(L+1)^2 + d(1/L+1)^2 * \\ (K/h)e(K/h) + f(h/K) + i(K/h)^2 + l(h/K)^2 * \\ (p/h)m(p/h) + n(h/p) + o(p/h)^2 + r(h/p)^2 * \\ (\mu)s(\mu) + t(1/\mu) + u(\mu)^2 + v(1/\mu)^2,$$

where the regression coefficients are $\{a, b, c, d, e, f, i, l, m, n, o, r, s, t, u, v\}$ and the fitting variables are $\{(L+1), (1/L+1), (L+1)^2, (1/L+1)^2$, and $(\theta/h), (h/\theta), (\theta/h)^2, (h/\theta)^2$ for $\theta = K, p$, and $\mu\}$. Each multiplicative model is logarithmically transformed to a linear expression for the regression runs. As Ehrhardt notes, "the resulting coefficients of determination ... [refer] to the proportion of variance explained for the logarithmically transformed variables [16]."

The backlog protection model, however, is slightly different from its cost counterparts. Ehrhardt comments that [16]:

The best approximation for backlog protection was obtained by departing from the [multiplicative] regression approach ... and using a theoretical result. When demand is continuously distributed, an optimal policy yields $\pi/(1+\pi)$ for backlog protection P [where $P = 1 - \text{stockout frequency}$]. Alternatively, if the demand distributions are discrete, $\pi/(1+\pi)$ is a lower bound on P for optimal policies.

where $\pi = p/h$. Thus, the backlog protection regression model is:

$$(5.2) \quad P = (c + a(L+1) + b(1/L+1) + c(L+1)^2 + d(1/L+1)^2 + \\ e(K/h) + f(h/K) + i(K/h)^2 + l(h/K)^2 + \\ m(p/h) + n(h/p) + o(p/h)^2 + r(h/p)^2 + \\ s(\mu) + t(1/\mu) + u(\mu)^2 + v(1/\mu)^2)/(1+\pi).$$

It should be noted that Equation 5.2 uses more fitting variables than the single penalty cost p variable used by Ehrhardt [16]; this was done to make the regression models more consistent.

5.2.2 Parameter Set

The regression models' fitted variables come from the operating characteristics of the 480 constrained optimal and constrained Power Approximation policies described in Sections 2.2.4 and 3.2.3, respectively. These variables are transformed as discussed above.

5.2.3 Evaluating Performance

The operating characteristic approximations are evaluated with two measures. The first measure is the coefficient of determination (r^2). Since these regressions are curve fits to mathematical functions, not statistically estimated functions, the r^2 value indicates how well the curve matches the operating characteristic values over the selected parameter range.

The second measure is the average percent deviation of the approximation's predicted value from the actual constrained optimum or constrained Power Approximation value. This deviation, calculated on a casewise basis and averaged over the 480 cases, is expressed on an absolute \bar{x}_a basis

$$(5.3) \quad \bar{x}_a = \frac{- \sum_{i=1}^n |(\hat{y}_i - y_i)/y_i|}{n} * 100,$$

where n is the number of observations, \hat{y}_i is the fitted operating characteristic value, and y_i is the actual value. Equation 5.3 emphasizes how helpful the approximation would be to an inventory manager trying to predict a given (s, S) inventory policy's operating characteristics.

The procedure for fitting a given approximation starts with running a stepwise regression for the log-linear model. Three of the stepwise models for each approximation are then selected for further evaluation based on their r^2 values. In this evaluation, each model's fitted values are compared to the characteristic's actual values and the average percent deviation calculated. An intermediate "best" model is then selected based on the model's size, r^2 value, and its average percent deviation \bar{x}_a . Based on this intermediate model, up to 5% of the parameter cases with the largest absolute percent deviations are eliminated as outliers (that is, poor fits). The stepwise regression is then repeated with the revised parameter set and three new models are chosen for further evaluation. Using the full 480 cases, a final "best" model for the performance approximation is chosen. This procedure is performed twice for each of the five performance characteristics, once each for the constrained optimal and Power Approximation policies.

A good example of this procedure is the development of the average total cost per period approximation for constrained optimal policies (T_{CO}/h). The procedure begins with running a stepwise regression using the log-linear version of equation (5.1) for the 480 constrained optimal policies. Three of the "better" regression models are applied to the 480 observations (variable number = 6, 5, and 4; $r^2 = 0.985$, 0.985, and 0.980; $\bar{x}_a = 5.4$, 5.4, and 6.6, respectively), and one "best" model is selected. The selected model is from step 5 of the regression

(5.4) $T_{CO}/h = 0.83592 * (L+1) \cdot 16628(L+1)$
 $* (K/h) \cdot 53012 + 1.6070h/K$
 $* (p/h) \cdot 12335 * (\mu) \cdot 46412.$

For this model, $r^2=0.985$, $\bar{X}_a=5.4\%$, and 57% of the 480 observations had absolute deviations $X_a < 5\%$. The 21 observations where $X_a > 0.15$ are deleted from the data set as "outliers" and the stepwise regression is repeated, again using the log-linear version of equation (5.1). The "best" (and final) model from this second regression is

$$(5.5) \quad T_{co}/h = 0.99143 * (L+1)^{.01374(L+1)} * (K/h)^{.53117+1.63582h/K} * (p/h)^{.00057(p/h)^2} * (\mu)^{.46624}.$$

In this case, $r^2 = 0.989$, $\bar{X}_a = 5.4\%$, and nearly 61% of the 480 observations had $X_a < 5\%$. This completes the development of the T_{co}/h approximation.

5.3 Evaluation of Results

Appendix III gives the final "best" models of each performance approximation for the constrained optimal and Power Approximation policies. Table 5.1 gives the performance criteria values for each of the final "best" models. On average, the approximations for backlog protection, total cost, replenishment cost, and holding cost tend to be unbiased, and their average absolute percent deviations do not exceed 12%. In contrast, the backlog cost approximations are poor. These approximations are biased towards overestimating the backlog cost and their average absolute percent deviations exceed 41%. While these findings are consistent with Ehrhardt [16], his approximations generally have larger r^2 values and smaller average absolute percent deviations.

The comparison with Ehrhardt's results is carried further in Table 5.2. Here the total cost approximations for the constrained optimal (co), constrained Power Approximation (cp), and unconstrained Power

Approximation (p) from Ehrhardt [16] are compared for single parameter variations. Note, that T_{co}/h is compared with constrained optimal policies, T_{cp}/h to the constrained Power Approximation, and T_p/h with the unconstrained Power Approximation.

Table 5.1

Operating Characteristic Approximation Performance for Constrained Optimal and Constrained Power Approximation Policies.

Operating Characteristic Approximation	r^2	\bar{X}_a	Percent of Cases	
			$X_a < 5\%$	$X_a < 10\%$
T_{co}/h	0.989	5.4*	60.6%	82.9%
T_{cp}/h	0.988	5.5	61.5	82.7
H_{co}/h	0.944	11.6	27.7	55.4
H_{cp}/h	0.957	9.6	34.6	62.1
R_{co}/h	0.989	7.4	47.5	75.0
R_{cp}/h	0.988	7.3	50.4	76.0
B_{co}/h	0.888	42.4	8.5	16.5
B_{cp}/h	0.872	41.4	10.2	16.7
P_{co}	0.997	1.8	96.3	100.0
P_{cp}	0.998	1.7	95.6	100.0

Legend:

T = Total cost

B = Backlog cost

H = Holding cost

P = Backlog protection

R = Replenishment cost

* X_a and \bar{X}_a values based on 480 cases.

Table 5.2

A Comparison of the Performance of Three Different Total Cost Approximations with Single Parameter Variations.
 Base Case: $K = 20$, $h = 0.5$, $p = 2$, $L = 2$, $\mu = 0.5$ (Poisson).

Changed Parameter	Percentage Above Actual Value		
	T_{co}/h	T_{cp}/h	T_p/h
Base Case	2.2	1.9	-40.7
K 3	2.9	5.8	-45.8
h 0.1	-2.1	-2.0	-46.0
0.3	0.3	-0.1	-41.5
0.7	3.4	3.3	-39.9
p $9*h$	-0.6	-0.1	-41.5
L 0	-6.4	-6.6	-43.2
4	4.2	3.8	-40.7
μ 0.1	-4.5	-4.3	-37.3
0.2	-0.4	-2.2	-40.4
0.3	1.3	-0.1	-40.4
0.4	2.0	1.8	-40.0
0.6	2.0	1.6	-40.5
0.7	1.6	1.2	-40.6
0.8	1.3	0.8	-40.5
0.9	0.8	0.3	-40.7
1.0	0.1	-0.6	-41.0

Legend: T_{co}/h = Constrained Optimal Total Cost Approx. T_{cp}/h = Constrained Power Approximation Total Cost Approx. T_p/h = Unconstrained Power Approximation Total Cost Approx.

Rarely do T_{co}/h and T_{cp}/h exceed 5% of their actual values.

Further, T_p/h consistently underestimates the total cost. It was derived for a parameter set having different demand distributions with larger mean demand values. These results confirm Ehrhardt's comment that performance approximations are not independent of their parameter sets [16]:

The fact that some of the approximations degrade when parameters are extrapolated is not surprising, since the expressions are numerical fits of curves that are empirically, not theoretically, motivated. The fitting approach we have described, however, can be used to derive different numerical approximations that would be accurate over other parameter ranges.

5.4 Applying the Approximations

This section shows how one of the operating characteristic approximations (T_{CO}/h) might be applied to three common inventory management problems: 1) gauging the sensitivity of average total cost to changing mean demand, 2) estimating the savings from consolidating warehouses, and 3) evaluating the benefit of shortening lead times.

5.4.1 The Sensitivity of Average Total Cost to Mean Demand

Inventory managers are often interested in how sensitive their inventory operations are to changes in the mean demand per period. For example, suppose the number of end items that use a particular inventory item is growing (or shrinking) and the manager wants to know how this will affect inventory costs. Equation (5.5) allows the manager to estimate the impact of changing mean demand.

Consider a base case where $K = \$20$, $h = \$0.50$, $p = \$2$, $L = 2$, and $\mu = 0.5$. We first provide a sensitivity analysis based on exact computations; we then repeat the analysis with the approximation. Under a constrained optimal (s, S) policy, the true expected total cost per period for the item is \$3.07. If demand increases 20% ($\mu = 0.6$), then the actual total cost per period increases to \$3.35. Likewise, if demand drops 20% ($\mu = 0.4$), then the actual total cost per period falls to \$2.77. Thus, a 20% increase or decrease in demand results in a +9% or -10% change in expected total costs per period, respectively.

Using equation (5.5) to achieve the sensitivity analysis, the expected total cost per period is approximated as \$3.14 for the base case. Further, if demand rises or falls by 20%, (5.5) estimates that the expected total cost per period is \$3.42 or \$2.83, respectively. Thus, a 20% increase or decrease in the mean demand per period results in a +9% or -10% change in total costs according to the approximation (5.5). Hence, the approximation for the expected total cost per period gives inventory managers the ability to quickly and accurately evaluate the cost impact of changing demand patterns.

5.4.2 Estimating Savings from Consolidating Warehouses

Frequently, inventory managers want to know whether it is cheaper to consolidate the stock for a low demand item at a single warehouse, rather than to distribute the stock among several warehouses. Consider again a constrained optimal (s, S) policy where $K = \$20$, $h = \$0.50$, $p = \$2$, and $L = 2$. If an item has a mean demand of 0.2 units per period, the actual expected total cost per period of stocking it at one warehouse is \$2.05. Stocking the item at two warehouses, therefore, would be \$4.10. If stocks are consolidated at one warehouse, so that $\mu = 0.4$ for that warehouse, the actual expected total cost is \$2.77 (exclusive of any additional shipping charges). Essentially, stocking the item at two warehouses is 32% more expensive than consolidating the inventory at one warehouse.

The approximation in (5.5) also allows an evaluation of whether to consolidate inventories. For the option of stocking the item at two warehouses ($\mu = 0.2$ units per period), equation (5.5) implies that the expected total cost per period is \$4.10. For the second option, where only one warehouse stocks the item ($\mu = 0.4$ units per period), the

approximate expected total cost per period is \$2.83. Thus the approximation indicates that there is a 31% reduction in costs due to consolidation. An inventory manager using the approximation (5.5) could quickly and relatively accurately compare the costs of stocking an item at a single warehouse versus multiple warehouses. Table 5.3 shows the cost data for the two, four, and eight warehouse cases. For each scenario, the approximation demonstrates its value as a quick and accurate predictor of actual costs.

5.4.3 Evaluating the Benefit of Shortening Lead Times

The last application shown for T_{co}/h helps determine the benefit of decreasing lead times. Inventory managers often encounter this problem when comparing several alternative modes of shipping, such as, air freight versus motor carrier.

Table 5.3

Applying the T_{co}/h Operating Characteristic Approximation for Determining Whether or Not to Consolidate Inventory at a Single Warehouse Location. ($K = \$20$, $h = \$0.5$, $p = \$2.0$, $L = 2$)

Number of Warehouses (μ)	Expected Total Cost per Period	
	Actual Cost	Approximation Cost
2 (0.2)	\$4.10	\$4.10
1 (0.4)	<u>2.77</u>	<u>2.83</u>
% Reduction	32%	31%
4 (0.2)	8.20	8.18
1 (0.8)	<u>3.86</u>	<u>3.91</u>
% Reduction	53%	52%
8 (0.1)	12.40	11.85
1 (0.8)	<u>3.86</u>	<u>3.91</u>
% Reduction	69%	67%

Consider a constrained optimal (s, S) policy for the case where $K = \$20$, $h = \$0.50$, $p = \$2$, and $\mu = 0.5$. We examine what is the impact of reducing L by 50%, such as, from 4 to 2. When $L = 2$ and 4, the actual expected total cost per period is \$3.07 and \$3.21, respectively. (Note that these amounts do not include transportation charges). Thus, in a low demand environment, reducing the delivery lag by 50% decreases costs by 4%. The approximation (5.5) suggests that when $L = 2$ and 4, the expected total cost per period is \$3.14 and \$3.35, respectively. This implies that decreasing lead times by 50% reduces total costs by 6%. Further, (5.5) suggests that reducing delivery lag by another 50% ($L = 1$) results in a cost of \$3.06, or only another 3% cost savings. Thus, (5.5) provides the inventory manager with valuable insights into the firm's opportunities for cost reductions. What is really important is the ease and accuracy of using operating characteristic approximations. Similar applications are feasible for the other approximations too.

5.5 Conclusions

For the 480 case parameter set, the constrained optimal and constrained Power Approximation operating characteristic approximations perform similarly. Although the accuracy of the approximations varies from characteristic to characteristic, the backlog protection, total cost, replenishment cost, and holding cost approximations perform well (within 12% average deviation). Unfortunately, the backlog cost approximations perform poorly and are not useful performance predictors for inventory managers. The regression derived operating characteristic approximations are best applied over their original parameter range; their performance degrades when they are used with parameters that are

unlike the set they were derived from. This chapter concludes the third and final research issue in this study. Chapter 6 summarizes the findings of this research, presents several implementation issues, and describes areas for future research.

CHAPTER 6

CONCLUSION

6.1 Research Summary

This research has examined the impact of low mean demand per period on the (s, S) periodic review inventory model. An unexpected result from applying the (s, S) model to low mean demand is that negative reorder points are commonly recommended. In other words, the firm does not reorder stock until a designated number of backorders are "on-the-books." For firms that cannot tolerate lengthy delays in ordering infrequently demanded but critical parts, the (s, S) model may be constrained to allow only nonnegative reorder points. While constraining reorder points increases the inventory service, it also increases total costs. Fortunately, the (s, S) model is relatively insensitive to misspecifying the mean demand per period. That is, using an estimated mean demand slightly higher or lower than the actual value does not increase total costs appreciably.

This research also showed that approximately optimal (s, S) inventory policies are feasible alternatives to computationally burdensome optimal (s, S) policies under low demand. In fact, there is very little cost degradation from using approximately optimal (s, S) policies rather than optimal (s, S) policies when demand is Poisson distributed. This is true whether the mean demand per period parameters are known or estimated from historical demand data.

Finally, useful approximations for (s, S) policy operating characteristics under low mean demand are easily developed. These approximations allow inventory practitioners to readily evaluate the impact of business decisions on their inventory costs and service.

6.2 Contributions of this Research

This research made two contributions to inventory theory. First, it examined the impact of low mean demand per period on the (s, S) periodic review inventory model, showing that negative reorder points are frequently recommended. This unexpected result had not been previously reported in the literature. Second, it confirmed Ehrhardt's [16] observation that approximations for (s, S) policy operating characteristics are dependent on the parameter set.

This research also made three contributions to inventory management. First, it demonstrated that the common practice of constraining reorder points to nonnegative values raises not only the service level, but also total system costs appreciably. Second, it showed that the Power and Analogy Approximations are practical alternatives to optimal (s, S) policies under low mean demand. Finally, the research demonstrated that useful approximations for (s, S) policy operating characteristics under low mean demand are easily developed.

6.3 Implementation Issues

This research provided insights into two major implementation issues confronting inventory managers. First, the results emphasize the importance of knowing the underlying nature of demand when using approximately optimal (s, S) periodic review inventory models. The performance of approximately optimal (s, S) policies under low mean demand varies with the distribution of demand, especially when the

demand parameters are statistically estimated from historical demand data. When demand is Poisson distributed the Analogy Approximation is the better choice. When there is more deviation in demand, however, the Power Approximation is the better choice. Thus, knowing the demand distribution of inventoried items permits managers to select an appropriate approximately optimal (s, S) inventory model.

Second, inventory managers should carefully examine the practical implications of their chosen inventory models over pertinent parameter ranges. For example, the negative reorder points which occur in some low demand cases could have serious ramifications to a service-oriented firm. The research shows that slightly modifying the inventory model's algorithm is an effective solution to this managerial dilemma.

6.4 Future Research

From this study, three main topics warrant future research: 1) the performance of constrained (s, S) policies given "real world" demand data, 2) alternative low demand environments, and 3) "mixed model" inventory systems.

6.4.1 Implementation Test

The notion of constrained (s, S) policies for infrequently demanded items needs to be tested using actual demand data from a firm using (s, S) policies. The key issue in this test is the comparison of total cost and service levels provided by the constrained and "real world" (s, S) policies. If the constrained (s, S) policies provide better inventory service for an acceptable increase in total costs, a trial implementation in an actual firm seems warranted.

6.4.2 Alternative Low Demand Environments

Under this heading are several different research issues. In 1985, Jacobs [26] suggested that robust, intentionally-biased estimators for the mean and variance of demand need to be evaluated under low mean demand. This research shows that approximately optimal (s, S) policies using the sample mean and variance result in large deviations from optimal total costs when the mean demand is small and its variance high. Perhaps robust, intentionally-biased estimators can improve the performance of approximately optimal (s, S) policies under these demand conditions.

Another research topic is the performance of (s, S) policies when the mean demand is low and nonstationary. For example, demand may follow a seasonal pattern or show a slight upward or downward trend. This environment is especially challenging to practitioners since their demand parameters are estimated from historical demand data. This may be another demand environment well-suited to intentionally-biased estimators.

A final demand-related topic is investigating the importance of lead time variability for low demand items. The inventory model used in this research assumed that the lead time was known. Since many periods can elapse between requirements for low demand items, however, is the (s, S) periodic review inventory model very sensitive to highly variable vendor lead times.

6.4.3 Mixed-Model Inventory Systems

The final research area should examine the usefulness of using an (s, S) inventory model that "selectively" constrains reorder points. How beneficial would it be for a firm to constrain the reorder points of

infrequently demanded items only if they are critical to the firm's operation? For example, a replacement part for an automated manufacturing system may warrant a constrained reorder point, whereas seldom used office equipment or cleaning supplies may not. This selective constraining is analogous to the ABC method of inventory control, where Class A and B items are the most important to the firm and therefore receive the most attention. The challenge to this "selectivity," however, is accurately identifying those vital items before they become critical.

APPENDIX I

480 OPTIMAL (s,S) POLICIES

480 OPTIMAL (S,S) POLICIES

OBS	H	U	B	K	L	P	R	S	T	S	HOLDING	PENALTY	STKOUTP
							OP	TKOBJ	TOTAL	SETUP			
1	0.1	0.10	20	0	0.40	-2	5	0.56	0.28	0.21	0.07	0.16	
2	0.1	0.10	20	0	0.90	-1	5	0.59	0.33	0.24	0.02	0.02	
3	0.1	0.10	20	2	0.40	-2	5	0.57	0.28	0.20	0.09	0.18	
4	0.1	0.10	20	2	0.90	-1	6	0.60	0.28	0.28	0.04	0.04	
5	0.1	0.10	20	4	0.40	-1	6	0.58	0.28	0.26	0.04	0.07	
6	0.1	0.10	20	4	0.90	-1	6	0.62	0.28	0.26	0.08	0.07	
7	0.1	0.30	20	0	1.20	-1	3	0.96	0.49	0.43	0.03	0.02	
8	0.1	0.30	20	0	2.70	-1	3	1.00	0.49	0.43	0.07	0.02	
9	0.1	0.30	20	2	1.20	-1	3	0.99	0.49	0.39	0.10	0.07	
10	0.1	0.30	20	2	2.70	-1	3	1.11	0.49	0.39	0.23	0.07	
11	0.1	0.30	20	4	1.20	-1	3	1.03	0.49	0.35	0.19	0.12	
12	0.1	0.30	20	4	2.70	0	4	1.19	0.49	0.61	0.08	0.03	
13	0.1	0.50	20	0	2.00	-1	2	1.20	0.66	0.48	0.07	0.03	
14	0.1	0.50	20	0	4.50	-1	2	1.29	0.66	0.48	0.15	0.03	
15	0.1	0.50	20	2	2.00	-1	2	1.30	0.66	0.41	0.23	0.10	
16	0.1	0.50	20	2	4.50	-1	3	1.53	0.49	0.65	0.38	0.07	
17	0.1	0.50	20	4	2.00	-1	3	1.39	0.49	0.59	0.31	0.12	
18	0.1	0.50	20	4	4.50	0	3	1.62	0.66	0.78	0.18	0.03	
19	0.1	0.70	20	0	2.80	-1	2	1.42	0.66	0.67	0.10	0.03	
20	0.1	0.70	20	0	6.30	-1	2	1.54	0.66	0.67	0.22	0.03	
21	0.1	0.70	20	2	2.80	-1	2	1.55	0.66	0.58	0.32	0.10	
22	0.1	0.70	20	2	6.30	-1	3	1.94	0.49	0.91	0.54	0.07	
23	0.1	0.70	20	4	2.80	-1	2	1.73	0.66	0.50	0.57	0.16	
24	0.1	0.70	20	4	6.30	0	3	2.00	0.66	1.09	0.26	0.03	
25	0.2	0.10	20	0	0.40	-2	8	0.80	0.40	0.35	0.06	0.12	
26	0.2	0.10	20	0	0.90	-1	8	0.85	0.44	0.39	0.02	0.02	
27	0.2	0.10	20	2	0.40	-2	8	0.81	0.40	0.32	0.09	0.16	
28	0.2	0.10	20	2	0.90	-1	9	0.87	0.40	0.40	0.07	0.06	
29	0.2	0.10	20	4	0.40	-2	9	0.82	0.36	0.34	0.13	0.18	
30	0.2	0.10	20	4	0.90	0	9	0.90	0.44	0.41	0.05	0.04	
31	0.2	0.30	20	0	1.20	-1	4	1.40	0.78	0.56	0.05	0.04	
32	0.2	0.30	20	0	2.70	-1	4	1.47	0.78	0.56	0.12	0.04	
33	0.2	0.30	20	2	1.20	-1	5	1.43	0.66	0.62	0.15	0.10	
34	0.2	0.30	20	2	2.70	-1	5	1.62	0.66	0.62	0.35	0.10	
35	0.2	0.30	20	4	1.20	-1	5	1.49	0.66	0.54	0.30	0.16	
36	0.2	0.30	20	4	2.70	0	6	1.66	0.66	0.79	0.22	0.06	
37	0.2	0.50	20	0	2.00	-1	3	1.78	0.98	0.70	0.11	0.05	
38	0.2	0.50	20	0	4.50	-1	3	1.91	0.98	0.70	0.24	0.05	
39	0.2	0.50	20	2	2.00	-1	4	1.89	0.78	0.80	0.31	0.12	
40	0.2	0.50	20	2	4.50	0	4	2.16	0.98	0.99	0.20	0.04	
41	0.2	0.50	20	4	2.00	-1	4	2.04	0.78	0.67	0.59	0.20	
42	0.2	0.50	20	4	4.50	0	5	2.29	0.78	1.07	0.44	0.07	
43	0.2	0.70	20	0	2.80	-1	3	2.10	0.98	0.97	0.15	0.05	
44	0.2	0.70	20	0	6.30	-1	3	2.29	0.98	0.97	0.34	0.05	
45	0.2	0.70	20	2	2.80	-1	3	2.30	0.98	0.79	0.53	0.15	
46	0.2	0.70	20	2	6.30	0	4	2.64	0.98	1.39	0.28	0.04	
47	0.2	0.70	20	4	2.80	0	4	2.48	0.98	1.16	0.34	0.09	
48	0.2	0.70	20	4	6.30	0	5	2.90	0.78	1.50	0.62	0.07	
49	0.3	0.10	20	0	0.40	-3	9	0.98	0.49	0.36	0.13	0.19	

480 OPTIMAL (S, S) POLICIES

480 OPTIMAL (s, S) POLICIES

480 OPTIMAL (s, S) POLICIES

O B S	H U	H K	L P	R O P	S T K O B J			T O L	S E T U P	H O L D I N G	P E N A L T Y	S T K O U T P	
					T	S	T						
148	0.7	0.10	20	2	0.90	0	18	1.64	0.76	0.77	0.11	0.07	
149	0.7	0.10	20	4	0.40	-1	18	1.54	0.72	0.62	0.20	0.18	
150	0.7	0.10	20	4	0.90	1	19	1.68	0.76	0.74	0.18	0.09	
151	0.7	0.30	20	0	1.20	-2	9	2.62	1.23	1.11	0.28	0.15	
152	0.7	0.30	20	0	2.70	-1	9	2.81	1.35	1.21	0.25	0.07	
153	0.7	0.30	20	2	1.20	-1	10	2.72	1.23	1.03	0.46	0.19	
154	0.7	0.30	20	2	2.70	1	11	3.00	1.35	1.40	0.26	0.06	
155	0.7	0.30	20	4	1.20	1	12	2.81	1.23	1.19	0.38	0.15	
156	0.7	0.30	20	4	2.70	2	13	3.14	1.23	1.45	0.46	0.09	
157	0.7	0.50	20	0	2.00	-2	7	3.43	1.50	1.37	0.57	0.18	
158	0.7	0.50	20	0	4.50	-1	7	3.72	1.68	1.53	0.51	0.08	
159	0.7	0.50	20	2	2.00	0	8	3.61	1.68	1.41	0.53	0.15	
160	0.7	0.50	20	2	4.50	1	9	4.04	1.68	1.83	0.53	0.07	
161	0.7	0.50	20	4	2.00	1	10	3.79	1.50	1.52	0.77	0.18	
162	0.7	0.50	20	4	4.50	3	11	4.32	1.68	2.13	0.51	0.06	
163	0.7	0.70	20	0	2.80	-1	6	4.07	1.90	1.80	0.36	0.10	
164	0.7	0.70	20	0	6.30	-1	6	4.52	1.90	1.80	0.81	0.10	
165	0.7	0.70	20	2	2.80	0	8	4.39	1.68	1.97	0.74	0.15	
166	0.7	0.70	20	2	6.30	1	8	4.97	1.90	2.23	0.84	0.08	
167	0.7	0.70	20	4	2.80	2	9	4.68	1.90	2.04	0.73	0.13	
168	0.7	0.70	20	4	6.30	3	10	5.36	1.90	2.64	0.81	0.07	
169	0.8	0.10	20	0	0.40	-4	16	1.61	0.78	0.64	0.19	0.19	
170	0.8	0.10	20	0	0.90	-2	17	1.71	0.82	0.75	0.14	0.09	
171	0.8	0.10	20	2	0.40	-2	18	1.63	0.78	0.67	0.17	0.17	
172	0.8	0.10	20	2	0.90	0	19	1.75	0.82	0.79	0.13	0.08	
173	0.8	0.10	20	4	0.40	-1	20	1.65	0.75	0.68	0.22	0.19	
174	0.8	0.10	20	4	0.90	2	21	1.79	0.82	0.83	0.13	0.07	
175	0.8	0.30	20	0	1.20	-2	9	2.80	1.40	1.09	0.31	0.16	
176	0.8	0.30	20	0	2.70	-1	10	3.01	1.40	1.34	0.27	0.07	
177	0.8	0.30	20	2	1.20	-1	11	2.91	1.29	1.11	0.51	0.19	
178	0.8	0.30	20	2	2.70	1	12	3.20	1.40	1.47	0.33	0.07	
179	0.8	0.30	20	4	1.20	1	13	3.00	1.29	1.23	0.48	0.17	
180	0.8	0.30	20	4	2.70	3	14	3.36	1.40	1.59	0.36	0.07	
181	0.8	0.50	20	0	2.00	-2	7	3.67	1.70	1.34	0.62	0.19	
182	0.8	0.50	20	0	4.50	-1	8	3.99	1.70	1.75	0.54	0.09	
183	0.8	0.50	20	2	2.00	0	9	3.86	1.70	1.54	0.61	0.16	
184	0.8	0.50	20	2	4.50	1	10	4.33	1.70	1.96	0.67	0.09	
185	0.8	0.50	20	4	2.00	2	11	4.05	1.70	1.74	0.61	0.14	
186	0.8	0.50	20	4	4.50	3	12	4.60	1.70	2.17	0.73	0.08	
187	0.8	0.70	20	0	2.80	-1	6	4.35	2.16	1.76	0.42	0.11	
188	0.8	0.70	20	0	6.30	0	7	4.82	2.16	2.39	0.27	0.03	
189	0.8	0.70	20	2	2.80	0	8	4.70	1.90	1.83	0.96	0.18	
190	0.8	0.70	20	2	6.30	1	9	5.35	1.90	2.41	1.04	0.10	
191	0.8	0.70	20	4	2.80	2	10	4.97	1.90	2.11	0.96	0.16	
192	0.8	0.70	20	4	6.30	3	11	5.75	1.90	2.70	1.14	0.09	
193	0.9	0.10	20	0	0.40	-4	17	1.70	0.84	0.68	0.19	0.18	
194	0.9	0.10	20	0	0.90	-2	18	1.82	0.88	0.80	0.14	0.09	
195	0.9	0.10	20	2	0.40	-2	19	1.72	0.84	0.70	0.19	0.17	
196	0.9	0.10	20	2	0.90	0	20	1.86	0.88	0.82	0.16	0.09	

480 OPTIMAL (s, S) POLICIES

480 OPTIMAL (s, S) POLICIES

480 OPTIMAL (s, S) POLICIES

O B S	B S U	H K L	P P J	S T K O B O		T T K O B O		S S U P		H O L D I N G		P E N A L T		S T K O U T P			
				T T K O B O		T T O B A J		S S U P		H O L D I N G		P E N A L T		S T K O U T P			
				S	U	H	K	L	P	S	U	T	P	G	I	S	T
295	0.3	0.30	3.00	0	1.20	-1	2	0.67	0.29	0.26	0.13	0.10					
296	0.3	0.30	3.00	0	2.70	-1	2	0.84	0.29	0.26	0.30	0.10					
297	0.3	0.30	3.00	2	1.20	0	3	0.82	0.29	0.38	0.15	0.10					
298	0.3	0.30	3.00	2	2.70	0	4	1.01	0.22	0.53	0.26	0.07					
299	0.3	0.30	3.00	4	1.20	0	4	0.94	0.22	0.40	0.32	0.17					
300	0.3	0.30	3.00	4	2.70	1	5	1.12	0.22	0.65	0.26	0.07					
301	0.3	0.50	3.00	0	2.00	-1	2	0.93	0.29	0.43	0.22	0.10					
302	0.3	0.50	3.00	0	4.50	0	2	1.14	0.42	0.63	0.09	0.02					
303	0.3	0.50	3.00	2	2.00	0	3	1.18	0.29	0.64	0.26	0.10					
304	0.3	0.50	3.00	2	4.50	1	3	1.46	0.42	0.84	0.20	0.04					
305	0.3	0.50	3.00	4	2.00	1	4	1.38	0.29	0.84	0.25	0.09					
306	0.3	0.50	3.00	4	4.50	1	4	1.69	0.29	0.84	0.57	0.09					
307	0.3	0.70	3.00	0	2.80	-1	1	1.15	0.42	0.28	0.45	0.14					
308	0.3	0.70	3.00	0	6.30	0	2	1.43	0.42	0.88	0.13	0.02					
309	0.3	0.70	3.00	2	2.80	0	2	1.51	0.42	0.58	0.52	0.14					
310	0.3	0.70	3.00	2	6.30	1	3	1.88	0.42	1.18	0.28	0.04					
311	0.3	0.70	3.00	4	2.80	1	3	1.77	0.42	0.85	0.50	0.12					
312	0.3	0.70	3.00	4	6.30	2	4	2.23	0.42	1.46	0.34	0.04					
313	0.4	0.10	3.00	0	0.40	-1	4	0.44	0.23	0.18	0.04	0.08					
314	0.4	0.10	3.00	0	0.90	-1	5	0.49	0.19	0.23	0.07	0.06					
315	0.4	0.10	3.00	2	0.40	-1	5	0.49	0.19	0.17	0.12	0.19					
316	0.4	0.10	3.00	2	0.90	0	6	0.55	0.19	0.25	0.10	0.08					
317	0.4	0.10	3.00	4	0.40	0	6	0.51	0.19	0.19	0.13	0.18					
318	0.4	0.10	3.00	4	0.90	1	7	0.59	0.19	0.27	0.13	0.09					
319	0.4	0.30	3.00	0	1.20	-1	2	0.80	0.38	0.25	0.18	0.13					
320	0.4	0.30	3.00	0	2.70	0	3	0.95	0.38	0.51	0.07	0.02					
321	0.4	0.30	3.00	2	1.20	0	4	0.95	0.29	0.46	0.21	0.12					
322	0.4	0.30	3.00	2	2.70	1	4	1.14	0.38	0.58	0.18	0.05					
323	0.4	0.30	3.00	4	1.20	1	5	1.07	0.29	0.53	0.25	0.13					
324	0.4	0.30	3.00	4	2.70	2	6	1.29	0.29	0.80	0.21	0.05					
325	0.4	0.50	3.00	0	2.00	-1	2	1.08	0.38	0.41	0.30	0.13					
326	0.4	0.50	3.00	0	4.50	0	2	1.30	0.55	0.59	0.16	0.03					
327	0.4	0.50	3.00	2	2.00	0	3	1.37	0.38	0.55	0.45	0.16					
328	0.4	0.50	3.00	2	4.50	1	4	1.65	0.38	0.97	0.31	0.05					
329	0.4	0.50	3.00	4	2.00	1	4	1.57	0.38	0.67	0.53	0.17					
330	0.4	0.50	3.00	4	4.50	2	5	1.91	0.38	1.08	0.45	0.07					
331	0.4	0.70	3.00	0	2.80	-1	2	1.37	0.38	0.57	0.42	0.13					
332	0.4	0.70	3.00	0	6.30	0	2	1.61	0.55	0.83	0.23	0.03					
333	0.4	0.70	3.00	2	2.80	0	3	1.76	0.38	0.76	0.63	0.16					
334	0.4	0.70	3.00	2	6.30	1	4	2.16	0.38	1.35	0.43	0.05					
335	0.4	0.70	3.00	4	2.80	1	4	2.05	0.38	0.93	0.74	0.17					
336	0.4	0.70	3.00	4	6.30	2	5	2.52	0.38	1.52	0.63	0.07					
337	0.5	0.10	3.00	0	0.40	-1	5	0.50	0.24	0.22	0.04	0.08					
338	0.5	0.10	3.00	0	0.90	-1	5	0.55	0.24	0.22	0.09	0.06					
339	0.5	0.10	3.00	2	0.40	0	6	0.54	0.24	0.23	0.07	0.12					
340	0.5	0.10	3.00	2	0.90	1	7	0.61	0.24	0.32	0.06	0.04					
341	0.5	0.10	3.00	4	0.40	1	7	0.57	0.24	0.24	0.10	0.14					
342	0.5	0.10	3.00	4	0.90	2	8	0.66	0.24	0.32	0.10	0.07					
343	0.5	0.30	3.00	0	1.20	-1	3	0.90	0.35	0.37	0.18	0.12					

AD-A186 151

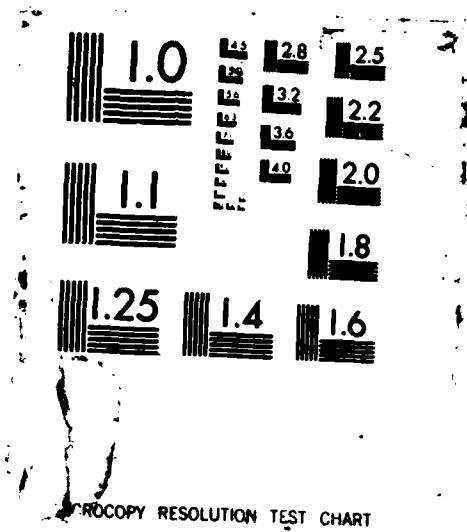
THE (S S) INVENTORY MODEL UNDER LOW DEMAND(U) AIR FORCE 2/2
INST OF TECH WRIGHT-PATTERSON AFB OH D K PETERSON 1987
AFIT/CI/NR-87-139D

UNCLASSIFIED

F/G 15/5

NL

FILE
NUMBER
151



480 OPTIMAL (s, S) POLICIES

O	S	R	S	T	S	H	P	S
B	U	O	K	O	E	O	E	T
S	U	B	R	T	S	D	S	K
		P	J	L	P	I	L	O
								U
344	0.5	0.30	3.00	0	2.70	0	3	1.05
345	0.5	0.30	3.00	2	1.20	0	4	1.08
346	0.5	0.30	3.00	2	2.70	1	5	1.26
347	0.5	0.30	3.00	4	1.20	1	6	1.21
348	0.5	0.30	3.00	4	2.70	2	7	1.45
349	0.5	0.50	3.00	0	2.00	-1	2	1.23
350	0.5	0.50	3.00	0	4.50	0	3	1.45
351	0.5	0.50	3.00	2	2.00	0	4	1.56
352	0.5	0.50	3.00	2	4.50	1	4	1.87
353	0.5	0.50	3.00	4	2.00	2	5	1.76
354	0.5	0.50	3.00	4	4.50	3	6	2.15
355	0.5	0.70	3.00	0	2.80	-1	2	1.54
356	0.5	0.70	3.00	0	6.30	0	2	1.80
357	0.5	0.70	3.00	2	2.80	1	4	2.00
358	0.5	0.70	3.00	2	6.30	1	4	2.43
359	0.5	0.70	3.00	4	2.80	2	5	2.28
360	0.5	0.70	3.00	4	6.30	3	6	2.82
361	0.6	0.10	3.00	0	0.40	-1	5	0.55
362	0.6	0.10	3.00	0	0.90	-1	6	0.61
363	0.6	0.10	3.00	2	0.40	0	7	0.59
364	0.6	0.10	3.00	2	0.90	1	8	0.67
365	0.6	0.10	3.00	4	0.40	1	8	0.63
366	0.6	0.10	3.00	4	0.90	2	9	0.72
367	0.6	0.30	3.00	0	1.20	-1	3	0.99
368	0.6	0.30	3.00	0	2.70	0	4	1.15
369	0.6	0.30	3.00	2	1.20	1	5	1.19
370	0.6	0.30	3.00	2	2.70	1	6	1.41
371	0.6	0.30	3.00	4	1.20	2	7	1.31
372	0.6	0.30	3.00	4	2.70	3	7	1.57
373	0.6	0.50	3.00	0	2.00	-1	3	1.38
374	0.6	0.50	3.00	0	4.50	0	3	1.57
375	0.6	0.50	3.00	2	2.00	1	4	1.69
376	0.6	0.50	3.00	2	4.50	2	5	2.04
377	0.6	0.50	3.00	4	2.00	2	6	1.92
378	0.6	0.50	3.00	4	4.50	3	7	2.34
379	0.6	0.70	3.00	0	2.80	-1	2	1.72
380	0.6	0.70	3.00	0	6.30	0	3	1.98
381	0.6	0.70	3.00	2	2.80	1	4	2.14
382	0.6	0.70	3.00	2	6.30	2	5	2.63
383	0.6	0.70	3.00	4	2.80	2	6	2.51
384	0.6	0.70	3.00	4	6.30	3	7	3.10
385	0.7	0.10	3.00	0	0.40	-1	6	0.59
386	0.7	0.10	3.00	0	0.90	-1	6	0.66
387	0.7	0.10	3.00	2	0.40	0	8	0.64
388	0.7	0.10	3.00	2	0.90	1	8	0.72
389	0.7	0.10	3.00	4	0.40	2	9	0.68
390	0.7	0.10	3.00	4	0.90	3	10	0.78
391	0.7	0.30	3.00	0	1.20	-1	3	1.09
392	0.7	0.30	3.00	0	2.70	0	4	1.23

480 OPTIMAL (s, S) POLICIES

O B S	B U H	K L P	R O P	S T K O J			T O L P			H O L D I N G			P E N A L T			S T K O U T P		
				S	T	K	T	O	L	S	E	N	A	L	T	Y	Y	T
393	0.7	0.30	3.00	2	1.20	1	6	1.28	0.39	0.67	0.22	0.11						
394	0.7	0.30	3.00	2	2.70	2	6	1.50	0.48	0.79	0.24	0.06						
395	0.7	0.30	3.00	4	1.20	2	7	1.42	0.39	0.60	0.43	0.18						
396	0.7	0.30	3.00	4	2.70	4	8	1.71	0.48	0.97	0.26	0.06						
397	0.7	0.50	3.00	0	2.00	-1	3	1.49	0.48	0.57	0.43	0.16						
398	0.7	0.50	3.00	0	4.50	0	3	1.70	0.63	0.74	0.33	0.06						
399	0.7	0.50	3.00	2	2.00	1	5	1.81	0.48	0.88	0.45	0.14						
400	0.7	0.50	3.00	2	4.50	2	6	2.18	0.48	1.31	0.39	0.06						
401	0.7	0.50	3.00	4	2.00	3	7	2.08	0.48	1.17	0.43	0.12						
402	0.7	0.50	3.00	4	4.50	4	8	2.53	0.48	1.61	0.43	0.06						
403	0.7	0.70	3.00	0	2.80	0	3	1.87	0.63	1.04	0.20	0.06						
404	0.7	0.70	3.00	0	6.30	0	3	2.13	0.63	1.04	0.46	0.06						
405	0.7	0.70	3.00	2	2.80	1	5	2.34	0.48	1.23	0.63	0.14						
406	0.7	0.70	3.00	2	6.30	2	5	2.83	0.63	1.49	0.71	0.07						
407	0.7	0.70	3.00	4	2.80	3	6	2.71	0.63	1.32	0.76	0.15						
408	0.7	0.70	3.00	4	6.30	4	7	3.31	0.63	1.92	0.77	0.07						
409	0.8	0.10	3.00	0	0.40	-1	6	0.64	0.32	0.25	0.06	0.11						
410	0.8	0.10	3.00	0	0.90	0	7	0.70	0.32	0.34	0.04	0.03						
411	0.8	0.10	3.00	2	0.40	0	8	0.68	0.29	0.26	0.14	0.18						
412	0.8	0.10	3.00	2	0.90	1	9	0.78	0.29	0.34	0.15	0.10						
413	0.8	0.10	3.00	4	0.40	2	10	0.72	0.29	0.30	0.14	0.16						
414	0.8	0.10	3.00	4	0.90	3	11	0.83	0.29	0.39	0.16	0.09						
415	0.8	0.30	3.00	0	1.20	-1	4	1.16	0.44	0.47	0.25	0.15						
416	0.8	0.30	3.00	0	2.70	0	4	1.32	0.55	0.58	0.20	0.06						
417	0.8	0.30	3.00	2	1.20	1	6	1.36	0.44	0.60	0.31	0.15						
418	0.8	0.30	3.00	2	2.70	2	7	1.60	0.44	0.86	0.29	0.07						
419	0.8	0.30	3.00	4	1.20	3	8	1.51	0.44	0.73	0.34	0.14						
420	0.8	0.30	3.00	4	2.70	4	9	1.80	0.44	0.99	0.37	0.08						
421	0.8	0.50	3.00	0	2.00	-1	3	1.61	0.55	0.55	0.51	0.18						
422	0.8	0.50	3.00	0	4.50	0	4	1.83	0.55	0.96	0.33	0.06						
423	0.8	0.50	3.00	2	2.00	1	5	1.96	0.55	0.78	0.63	0.18						
424	0.8	0.50	3.00	2	4.50	2	6	2.34	0.55	1.19	0.60	0.08						
425	0.8	0.50	3.00	4	2.00	3	7	2.22	0.55	0.99	0.68	0.17						
426	0.8	0.50	3.00	4	4.50	4	8	2.70	0.55	1.41	0.75	0.09						
427	0.8	0.70	3.00	0	2.80	0	3	1.97	0.71	1.00	0.26	0.07						
428	0.8	0.70	3.00	0	6.30	0	3	2.30	0.71	1.00	0.59	0.07						
429	0.8	0.70	3.00	2	2.80	1	5	2.52	0.55	1.09	0.88	0.18						
430	0.8	0.70	3.00	2	6.30	2	6	3.06	0.55	1.67	0.84	0.08						
431	0.8	0.72	3.00	4	2.80	3	7	2.89	0.55	1.39	0.95	0.17						
432	0.8	0.70	3.00	4	6.30	5	8	3.57	0.71	2.26	0.60	0.06						
433	0.9	0.10	3.00	0	0.40	-1	7	0.68	0.32	0.29	0.06	0.11						
434	0.9	0.10	3.00	0	0.90	0	7	0.75	0.36	0.33	0.05	0.04						
435	0.9	0.10	3.00	2	0.40	0	9	0.73	0.29	0.29	0.15	0.19						
436	0.9	0.10	3.00	2	0.90	2	10	0.82	0.32	0.41	0.09	0.06						
437	0.9	0.10	3.00	4	0.40	2	11	0.77	0.29	0.31	0.17	0.18						
438	0.9	0.10	3.00	4	0.90	4	12	0.88	0.32	0.43	0.13	0.07						
439	0.9	0.30	3.00	0	1.20	-1	4	1.24	0.50	0.46	0.29	0.17						
440	0.9	0.30	3.00	0	2.70	0	5	1.40	0.50	0.71	0.20	0.06						
441	0.9	0.30	3.00	2	1.20	1	7	1.45	0.42	0.68	0.35	0.16						

480 OPTIMAL (S,S) POLICIES

O B S	H U P	E K L	R O P	T O J	T S L	Z U P	D I G	I N G	L T Y	H E G	P E G	S T Y	
442	0.9	0.30	3.00	2	2.70	2	7	1.71	0.50	0.79	0.43	0.09	
443	0.9	0.30	3.00	4	1.20	3	9	1.61	0.42	0.76	0.43	0.17	
444	0.9	0.30	3.00	4	2.70	5	10	1.92	0.50	1.14	0.29	0.06	
445	0.9	0.50	3.00	0	2.00	0	3	1.70	0.78	0.68	0.23	0.09	
446	0.9	0.50	3.00	0	4.50	0	4	1.95	0.61	0.93	0.41	0.07	
447	0.9	0.50	3.00	2	2.00	2	6	2.07	0.61	1.08	0.39	0.11	
448	0.9	0.50	3.00	2	4.50	3	7	2.49	0.61	1.52	0.36	0.05	
449	0.9	0.50	3.00	4	2.00	4	8	2.36	0.61	1.22	0.53	0.14	
450	0.9	0.50	3.00	4	4.50	5	9	2.84	0.61	1.65	0.59	0.07	
451	0.9	0.70	3.00	0	2.80	0	3	2.07	0.78	0.96	0.33	0.09	
452	0.9	0.70	3.00	0	6.30	0	3	2.48	0.78	0.96	0.74	0.09	
453	0.9	0.70	3.00	2	2.80	2	5	2.66	0.78	1.19	0.69	0.14	
454	0.9	0.70	3.00	2	6.30	3	6	3.21	0.78	1.79	0.64	0.06	
455	0.9	0.70	3.00	4	2.80	4	8	3.06	0.61	1.70	0.75	0.14	
456	0.9	0.70	3.00	4	6.30	5	9	3.74	0.61	2.31	0.82	0.07	
457	1.0	0.10	3.00	0	0.40	-1	7	0.71	0.35	0.29	0.07	0.12	
458	1.0	0.10	3.00	0	0.90	0	8	0.78	0.35	0.38	0.05	0.04	
459	1.0	0.10	3.00	2	0.40	1	10	0.77	0.32	0.35	0.10	0.13	
460	1.0	0.10	3.00	2	0.90	2	11	0.86	0.32	0.43	0.11	0.07	
461	1.0	0.10	3.00	4	0.40	3	12	0.81	0.32	0.36	0.14	0.15	
462	1.0	0.10	3.00	4	0.90	4	13	0.93	0.32	0.44	0.17	0.09	
463	1.0	0.30	3.00	0	1.20	-1	4	1.31	0.55	0.44	0.33	0.18	
464	1.0	0.30	3.00	0	2.70	0	5	1.48	0.55	0.69	0.25	0.07	
465	1.0	0.30	3.00	2	1.20	2	7	1.53	0.55	0.72	0.26	0.12	
466	1.0	0.30	3.00	2	2.70	3	8	1.79	0.55	0.99	0.26	0.06	
467	1.0	0.30	3.00	4	1.20	4	10	1.70	0.46	0.90	0.34	0.13	
468	1.0	0.30	3.00	4	2.70	5	11	2.02	0.46	1.16	0.40	0.08	
469	1.0	0.50	3.00	0	2.00	0	4	1.79	0.67	0.90	0.22	0.08	
470	1.0	0.50	3.00	0	4.50	0	4	2.06	0.67	0.90	0.50	0.08	
471	1.0	0.50	3.00	2	2.00	2	6	2.17	0.67	0.98	0.53	0.15	
472	1.0	0.50	3.00	2	4.50	3	7	2.60	0.67	1.40	0.53	0.07	
473	1.0	0.50	3.00	4	2.00	4	9	2.47	0.55	1.26	0.66	0.16	
474	1.0	0.50	3.00	4	4.50	5	10	3.01	0.55	1.69	0.78	0.09	
475	1.0	0.70	3.00	0	2.80	0	3	2.17	0.86	0.92	0.40	0.10	
476	1.0	0.70	3.00	0	6.30	0	4	2.62	0.67	1.26	0.70	0.08	
477	1.0	0.70	3.00	2	2.80	2	6	2.78	0.67	1.37	0.74	0.15	
478	1.0	0.70	3.00	2	6.30	3	7	3.37	0.67	1.96	0.74	0.07	
479	1.0	0.70	3.00	4	2.80	4	9	3.24	0.55	1.77	0.93	0.16	
480	1.0	0.70	3.00	4	6.30	6	10	-3.97	0.67	2.65	0.65	0.06	

APPENDIX II
SUPPLEMENTARY SIMULATION RESULTS

Comparison of Constrained Optimal with the Constrained Power and Analogy Approximations' Operating Characteristics Under Known and Estimated Means - 96 Item System - Negative Binomial

Policy	Average					Stockout Frequency	
	Costs				Penalty		
	Total	Setup	Hold				
Optimal (kn)	\$3.23	\$.53	\$1.29	\$1.42		0.09	
Power (kn)	3.59	.39	2.31	.89		0.05	
Power (est)	3.90	.62	2.20	1.08		0.07	
Analogy (kn)	3.71	.37	2.49	.85		0.05	
Analogy (est)	4.02	.58	2.48	.96		0.06	

Comparison of Operating Characteristics for Two Constrained Approximation Policies to the Optimum When Mean Demand is Known and Estimated - 96 Item System - Negative Binomial

Policy	Average Percent Deviation					Stockout Frequency	
	Costs				Penalty		
	Total	Setup	Hold				
Power (kn)	11.3	-25.7	79.9	-37.2		-40.6	
Power (est)	20.8	18.0	71.2	-23.9		-20.5	
Analogy (kn)	14.9	-29.7	93.5	-39.9		-45.4	
Analogy (est)	24.6	9.3	93.2	-31.9		-29.4	

APPENDIX III

**OPERATING CHARACTERISTIC APPROXIMATION MODELS FOR THE
CONSTRAINED OPTIMAL AND POWER APPROXIMATION POLICIES**

1. Average Total Cost per Period.

$$T_{co}/h = 0.99143 * (L+1) \cdot 01374(L+1) * (K/h) \cdot 53117 + 1.63582h/K * (p/h) \cdot 00057(p/h)^2 * (\mu) \cdot 46624 .$$

$$T_{cp}/h = 0.90949 * (L+1) \cdot 01361(L+1) * (K/h) \cdot 54571 + 1.92091h/K * (p/h) \cdot 00065(p/h)^2 * (\mu) \cdot 46324 .$$

2. Average Holding Cost per Period.

$$H_{co}/h = 0.34469 * (L+1) \cdot 00194(L+1)^2 * (K/h) \cdot 57422 + 2.57211h/K * (p/h) \cdot 00108(p/h)^2 * (\mu) \cdot 46047 - 0.00648/\mu .$$

$$H_{cp}/h = 0.27667 * (L+1) \cdot 01822(L+1) * (K/h) \cdot 59683 + 3.10752h/K * (p/h) \cdot 00168(p/h)^2 * (\mu) \cdot 32225 + .11538 .$$

3. Average Replenishment Cost per Period.

$$R_{co}/h = 0.57824 * (L+1) \cdot 14488 + 0.00683(L+1) * (K/h) \cdot 54598 - 1.6358(h/K)^2 * (\mu) \cdot 46979 .$$

$$R_{cp}/h = 0.5418 * (L+1) \cdot 12971 * (K/h) \cdot 53445 * (p/h) \cdot 47486h/p * (\mu) \cdot 55257 .$$

4. Average Backlog Cost per Period.

$$B_{co}/h = 1.96273 * (L+1) \cdot 1.67008 - 0.08413(L+1) * (K/h) \cdot .20163 * (p/h) \cdot 4.18764h/p * (\mu) \cdot 1.03536 .$$

$$B_{cp}/h = 0.35841 * (L+1) \cdot 1.44113 - 0.0678(L+1) * (K/h) \cdot .00024K/h * (p/h) \cdot .83758h/p * (\mu) \cdot 91079 + .01507/\mu .$$

5. Backlog Protection.

$$P_{co} = (13.1688 + .57094/(L+1) + .0009K/h - 35.50327h/p - .28027\mu + .0284/\mu)/(1 + \pi).$$

$$P_{cp} = (13.78229 + .45953/(L+1) - 36.24002h/p - .94459(\mu) + .4299(\mu)^2)/(1 + \pi).$$

REFERENCES

1. Arrow, Kenneth J., Theodore Harris, and Jacob Marschak, "Optimal Inventory Policy," Econometrica, Vol. 19, No. 3 (1951), pp. 250-272.
2. Arrow, Kenneth J., Samuel Karlin and Herbert Scarf, (eds.) "The Nature and Structure of Inventory Problems," Chapter 2 in Studies in the Mathematical Theory of Inventory and Production, Stanford, CA: Stanford University Press, 1958.
3. Bellman, R., I. Glicksberg, R. Gross, "On the Optimal Inventory Equation," Management Science, Vol. 2, (1955), pp. 83-104.
4. Blazer, Douglas J., "An Inventory Model that Specially Handles Extreme Value Demand," Ph.D. dissertation, School of Business Administration, University of North Carolina, Chapel Hill, N.C., 1983.
5. Brown, Bernice B., Characteristics of Demand for Aircraft Spare Parts, The RAND Corporation R-292, July 1956.
6. Davis, D. J., "An Analysis of Some Failure Data," Journal of the American Statistical Association, Vol. 47, No. 256 (June 1952), pp. 113-150.
7. Dvoretzky, A., J. Kiefer, J. Wolfowitz, "The Inventory Problem: I - Case of Known Distributions of Demand," Econometrica, Vol. 20, No. 2 (1952), pp. 187-222.
8. Dvoretzky, A., J. Kiefer, J. Wolfowitz, "The Inventory Problem: II - Case of Unknown Distributions of Demand," Econometrica, Vol. 20, No. 3 (1952), pp. 450-466.
9. Dvoretzky, A., J. Kiefer, J. Wolfowitz, "On the Optimal Character of the (s,S) Policy in Inventory Theory," Econometrica, Vol. 21, (1953), pp. 586-596.
10. Ehrhardt, Richard, "Power Approximation: Inventory Policies Based on Limited Demand Information," ONR and ARO Technical Report 7, New Haven, CT: School of Organization and Management, Yale University, 1976.
11. Ehrhardt, Richard, "The Power Approximation for Computing (s,S) Inventory Policies," Management Science, Vol. 25, No. 8 (1979), pp. 777-786.

12. Ehrhardt, R. and G. Kastner, "An Empirical Comparison of Two Approximately Optimal (s, S) Inventory Policies," ONR Technical Report 16, School of Business Administration, University of North Carolina at Chapel Hill, Chapel Hill, N.C., December 1980.
13. Ehrhardt, R., and H. M. Wagner, "Inventory Models and Practice," in Advanced Techniques in the Practice of Operations Research, H. J. Greenberg, F. H. Murphrey, and S. H. Shaw (eds.), New York: Elsevier, 1982, pp. 250-332.
14. Ehrhardt, Richard, "Variance Reduction Techniques for (s, S) -Type Inventory Simulations," in Simulation in Inventory and Production Control, H. Bekiroglu (ed.), La Jolla: Simulation Counsels Inc., 1983.
15. Ehrhardt, Richard and Charles Mosier, "A Revision of the Power Approximation for Computing (s, S) Policies," Management Science, Vol. 30, No. 5 (May 1984), pp. 618-622.
16. Ehrhardt, Richard, "Easily Computed Approximations for (s, S) Inventory System Operating Characteristics," Naval Research Logistics Quarterly, Vol. 32, 1985, pp. 347-359.
17. Estey, A. S. and R. L. Kaufman, "Multi-Item Inventory System Policies Using Statistical Estimates: Negative Binomial Demands (Variance/Mean = 9)," ONR and ARO Technical Report 3, School of Organization and Management, Yale University, New Haven, CT, 1975.
18. Federgruen, Awi, and P. Zipkin, "An Efficient Algorithm for Computing Optimal (s, S) Policies," Operations Research, Vol. 32, No. 6 (Nov-Dec 1984), pp. 1268-1285.
19. Feeney, G. E., J. W. Petersen, and C. C. Sherbrooke, An Aggregate Base Stockage Policy for Recoverable Spare Parts, The RAND Corporation Memorandum, RM-3644-PR, June 1963.
20. Frazza, Sally, Test of Poisson Failure Assumption, Philadelphia, PA: US Army Material Systems Analysis Activity, September 1982.
21. Freeland, James R., and Evan L. Porteus, "Evaluating the Effectiveness of a New Method for Computing Approximately Optimal (s, S) Inventory Policies," Operations Research, Vol. 28, No. 2 (March-April 1980), pp. 353-364.
22. Haber, Sheldon E., and Rosedith Sitgreaves, "A Methodology for Estimating Expected Usage of Repair Parts With Application to Parts With No Usage History," Naval Research Logistics Quarterly, Vol. 12, No. 4 (December 1970), pp. 535-546.
23. Haber, Sheldon E., and Rosedith Sitgreaves, "A Unified Model for Demand Prediction in the Context of Provisioning and Replenishment," Naval Research Logistics Quarterly, Vol. 19, No. 1, (March 1972), pp. 29-32.

24. Heyvaert, A. C. and A. Hurt, "Inventory Management of Slow-Moving Parts," Operations Research, Vol. 4, No. 5 (October 1956), pp. 572-580.
25. Iglehart, Donald L., "Dynamic Programming and Stationary Analysis of Inventory Problems," Chapter 1 in Multistage Inventory Models and Techniques, H. E. Scarf, D. M. Gilford, and M. W. Shelly (eds.), Stanford, CA: Stanford University Press, 1963.
26. Jacobs, Raymond A., "Statistical Forecasting of Demand for Inventory Control Systems," Ph.D. dissertation, School of Business Administration, University of North Carolina, Chapel Hill, N.C., 1985.
27. Johnson, E., "On (s,S) Policies," Management Science, Vol. 15, (1968), pp. 80-101.
28. Kaufman, R., "Computer Programs for (s,S) Policies Under Independent or Filtered Demands," ONR and ARO Technical Report 5, New Haven, CT: School of Organization and Management, Yale University, 1976.
29. Kaufman, R., "(s,S) Inventory Policies in a Nonstationary Demand Environment," ONR and ARO Technical Report 11, School of Organization and Management, Yale University, New Haven, CT, 1977.
30. Klincewicz, John G., "Biased Variance Estimators for Statistical Inventory Policies," ONR and ARO Technical Report 8, School of Organization and Management, Yale University, New Haven, CT, 1976.
31. Klincewicz, John G., "Inventory Control Using Statistical Estimates: The Power Approximation and Sporadic Demands (Variance/ Mean = 9)," ONR and ARO Technical Report 9, Chapel Hill, NC: School of Business Administration, University of North Carolina at Chapel Hill, 1976.
32. MacCormick, A., "Statistical Problems in Inventory Control," ONR and ARO Technical Report 2, School of Organization and Management, Yale University, New Haven, CT, 1974.
33. Mosier, Charles, "Revised (s,S) Power Approximation," ONR Technical Report 18, Chapel Hill, NC: School of Business Administration, University of North Carolina at Chapel Hill, 1981.
34. Naddor, Eliezer, "Optimal and Heuristic Decisions in Single- and Multi-Item Inventory Systems," Management Science, Vol. 21, No. 11 (July 1975), pp. 1234-1249.
35. Porteus, Evan L., "Numerical Comparisons of Inventory Policies for Periodic Review Systems," Operations Research, Vol. 33, No. 1 (Jan-Feb 1985), pp. 134-152.

36. Roberts, D., "Approximations to Optimal Policies in a Dynamic Inventory Model," Studies in Applied Probability and Management Science, K. Arrow, S. Karlin, and H. Scarf (eds.), Stanford, CA: Stanford University Press, 1962.
37. Scarf, Herbert E., "Bayes Solutions of the Statistical Inventory Problem," Annals of Mathematical Statistics, Vol. 30, No. 2 (1959), pp. 490-508.
38. Scarf, Herbert E., "Some Remarks on Bayes Solutions to the Inventory Problem," Naval Research Logistics Quarterly, Vol. 7, No. 4 (1960), pp. 591-596.
39. Scarf, Herbert E., "The Optimality of (S,s) Policies in the Dynamic Inventory Problem," Chapter 13 in Mathematical Models in the Social Sciences, K. J. Arrow, S. Karlin and P. Suppes (eds.), Stanford, CA: Stanford University Press, 1960.
40. Scarf, Herbert E., "A Survey of Analytic Techniques in Inventory Theory," Chapter 7 in Multistage Inventory Models and Techniques, H. Scarf, D. Gilford, and M. Shelly (eds.), Stanford, CA: Stanford University Press, 1963.
41. Schaefer, Margaret K., "A Multi-Item Maintenance Center Inventory Model for Low-Demand Reparable Items," Management Science, Vol. 29, No. 9, (September 1983), pp. 1062-1068.
42. Silver, E. A., "Operations Research in Inventory Management: A Review and Critique," Operations Research, Vol. 29, No. 4 (1981), pp. 628-645.
43. Timjs, H. C., Analysis of (s,S) Inventory Models, Amsterdam: Mathematisch Centrum, 1972.
44. U.S. Department of the Air Force, Stockage Policy Handbook for Supply Officers, AFLMC/LGS, Gunter AFS, AL, 1985.
45. Veinott, A. F., and H. M. Wagner, "Computing Optimal (s,S) Inventory Policies," Management Science, Vol. 11, No. 5 (1965), pp. 525-552.
46. Wagner, Harvey M., M. O'Hagan, and B. Lundh, "An Empirical Study of Exactly and Approximately Optimal Inventory Policies," Management Science, Vol. 11, No. 7 (May 1965), pp. 690-723.
47. Wagner, Harvey M., Principles of Operations Research, Englewood Cliffs, NJ: Prentice-Hall Inc., 1969.
48. Wagner, Harvey M., "Research Portfolio for Inventory Management and Production Planning Systems," Operations Research, Vol. 28, No. 3 (May-June 1980), pp. 445-474.
49. Whitin, T. M. and J. W. T. Youngs, "A Method for Calculating Optimal Inventory Levels and Delivery Time," Naval Research Logistics Quarterly, Vol. 2, No. 3, (September 1955), pp. 157-173.

END

DATE

FILMED

DEC.

1987